

UBUNTU
user

UBUNTU 9.04

JAUNTY JACKALOPE
DOUBLE-SIDED DVD

32-BIT & 64-BIT
VERSIONS

UBUNTU



user

EXPLORING THE WORLD OF UBUNTU

JUMP TO JAUNTY

Ubuntu 9.04 "Jaunty Jackalope"
Faster boot and better tools

Run Windows programs on Ubuntu

Security tricks

Virtualization with VirtualBox

10 COOL GAMES!

Jono Bacon: Welcome to
the diverse Ubuntu community

Exclusive interview with
Ubuntu kernel manager
Pete Graner



DISCOVERY GUIDE

Our handy guide helps you
get started with Ubuntu
"Linux for human beings" p61



WWW.UBUNTU-USER.COM

£7.99
A\$19.99
Issue 01





Google

FREE DVD! Google in Ubuntu

Google Tools

LINUX MAGAZINE



GOOGLE SPECIAL

Our custom DVD includes Google's online tools in an easy desktop menu. The Live boot option lets you try the DVD without disturbing your current configuration.



Fall 2009

Google

Get Productive

Collaborate from anywhere with Google's office tools

Living Online

AVAILABLE AT YOUR NEWSSTAND!

OR ORDER ONLINE AT WWW.LINUX-MAGAZINE.COM/SPECIAL

FREE TOOLS!

- Manage online shopping from a single web interface
- Create custom illustrations
- Build your own blog
- Improve your site's search rank



Catch the Wave!



Will Google's powerful Wave app change the Internet?

Special Edition 05



9 771757 636002

05

£ 7.99



Your Ubuntu 9.04 DVD

This Ubuntu “Jaunty Jackalope” DVD includes 32-bit, 64-bit, and Live versions of Ubuntu 9.04.

If you are planning on installing Ubuntu, use the 32-bit version unless you are sure your hardware supports 64-bit operations (see your hardware vendor documentation).

The Live version lets you boot your system from the DVD without disturbing the contents of your hard drive. You can use the Live boot option to check the compatibility of your system before you commit to a full installation. Be sure to back up any essential data before you install.

This DVD supports most modern Intel, or equivalent, PCs. To install Ubuntu, insert the the disc and restart your system. Note that this DVD is a double-sided disc. This DVD is labeled like a normal DVD – the label you see as you place the disc in the tray applies to the data on the opposite side.

If your system will not boot to the DVD, you might need to reconfigure the boot order in your computer’s BIOS menu. See the the Discovery Guide (starting on page 61) for more on installing and exploring Ubuntu.

Exploring the world of Ubuntu

JUMP TO JAUNTY

Welcome to *Ubuntu User Magazine* – your independent guide to the Ubuntu experience. Inside you’ll find helpful tools, great ideas, and plenty of good advice on configuring and troubleshooting your Ubuntu system.

If you’re new to Ubuntu, this issue provides you with all you need to get started. The DVD opposite this page includes 32-bit and 64-bit versions of Ubuntu 9.04 “Jaunty Jackalope,” and our exclusive Discovery Guide (starting on page 61) shows you how to get your Ubuntu system up and running.

Ubuntu veterans will also find much to learn in this issue, with insightful articles from experts such as Marcel Gagné and Klaus Knopper, as well as tips from Ubuntu insiders Mike Basinger and Jono Bacon.

Our lead-off articles examine graphics in Ubuntu and how to run Windows as a virtual system – but that’s just the beginning. Later articles explore popular Ubuntu

games, show how to set up a firewall, and describe how to build custom software packages.


We also tour some exciting desktop tools that are freely available for all Ubuntu users, such as the FreeMind mind mapper and the Celtx authoring application.

If you’re an Ubuntu user, we want to hear from you. Visit our website to learn more about what’s coming in future issues of *Ubuntu User*. ubuntu-user.com



Additional Resources

- [1] Ubuntu: <http://www.ubuntu.com/>
- [2] Ubuntu Forums: <http://ubuntuforums.org/>
- [3] Ubuntu wiki: <https://wiki.ubuntu.com/>
- [4] Ubuntu Support overview: <http://www.ubuntu.com/support>



Jaunty Jackalope: Ubuntu User is your companion for exploring the world of Ubuntu. This issue celebrates the arrival of

Ubuntu 9.04 "Jaunty Jackalope." Read on for expert tools and techniques at the center of the Ubuntu experience. **6**

10 Great Games: All work and no play? We show you how to get your game on with 10 cool game applications. **16**

NEWS

- 6 New Features**
Jaunty Jackalope offers faster boot times, notifications, desktop improvements, and more.
- 10 Who We All Are**
Jono Bacon welcomes you to the diverse universe of Ubuntu.

FEATURES

- 11 Virtualization**
We show you how to operate virtual environments from your Ubuntu desktop.
- 16 10 Great Games**
Relax! You work too hard! We tour some top picks for great gaming in Ubuntu.
- 20 Interview: Pete Graner**
Ubuntu's Kernel Manager talks about development and community.



SERVICE

- 3 Ubuntu 9.04 DVD**
- 97 About Us**
- 98 Preview**

KNOW-HOW

- 22 Graphics Cards**
Learn how to configure accelerated graphics cards in Ubuntu.

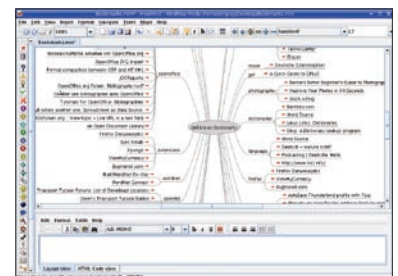


- 26 UFW Firewall**
Secure your system and keep intruders out with a personal firewall.
- 30 Wine**
Maybe you don't have to give up Windows programs after all. Learn how to run Windows applications on Ubuntu with Wine.
- 36 Building a Website**
Learn about some tools for building and publishing web pages.
- 42 Integrating Debian Packages**
We show you how to create custom Ubuntu software packages using Debian package format.
- 46 Answer Man**
Get answers to questions about playing DVDs, setting up email, installing from a USB stick, fixing the Ubuntu boot menu, and more.



DESKTOP


- 48 Gnome Do**
Search engines meet the application launcher with the handy Gnome Do tool.
- 50 FreeMind**
Looking for inspiration? We show you how to visualize your best ideas with this useful mind-mapping app.



- 53 Celtx Authoring Software**
This easy authoring tool is the perfect companion for a dark and stormy night. We show you how to compose comics, movie scripts, and novels with Celtx.
- 56 Synergy**
Control multiple PCs with a single keyboard and mouse, even across operating system boundaries.
- 58 Environment Variables**
We show you some terminal tricks for working with environment variables.



UFW Firewall: Don't miss the chance to lock out intruders with Ubuntu's built-in firewall application. 26



Celtx: Got an idea for a novel or play? Get your creative juices going with this handy authoring tool. 53

Double-sided DVD
Includes 32-bit, 64-bit, and Live versions



UBUNTU 9.04 JAUNTY JACKALOPPE

- Full Ubuntu 9.04 on a double-sided DVD
- Hard-disk install + Live versions
- 10GB Ubuntu DVD
- Includes 32-bit, 64-bit, and Live versions

See p3 for full details

Discovery Guide



Peek inside for a practical introduction to Ubuntu, including tips on installing, configuring, and exploring your new Ubuntu system.

Installing Ubuntu	62	Configuring Email.....	84
Exploring.....	64	Taking Notes	88
Networking	67	OpenOffice	90
Security.....	71		
Printing	74		
Multimedia	76		
Burning CDs.....	78		
Managing Photos.....	80		
Installing Software	82		



What's new in Ubuntu 9.04 Jaunty Jackalope

ANNOUNCING THE RABBIT

Speed was the theme for the latest Ubuntu development effort, so what could be a better symbol than a mythical creature that is half antelope and half jackrabbit? Jaunty Jackalope offers faster boot times and lots of other improvements. **BY KRISTIAN KISSLING, NICK ALI, AND JOE CASAD**


On April 23, 2009, the final release of Ubuntu 9.04 rolls out to the public [1]. Download links light up around the world as the version known as “Jaunty Jackalope” hops to center stage. The update has been in the works for months, and Ubuntu insiders are already working with pre-release versions – alphas, betas, release candidates – but the moment of the final release is still a special time for any software development team.

Sister distributions Kubuntu and Xubuntu appear nearly simultaneously, and other projects such as the Edubuntu, Ubuntu Studio, and the Netbook Remix edition make their way to the public. Another Ubuntu is out in the world.

If you follow the fortunes of Ubuntu, or if you are new to this stylish Linux system that has become a favorite with users around the world, you might be wondering what's new in the latest release. Following is a brief look at some of the newest features. But beyond all the tools and options, you'll also find new art, better performance, and more comprehensive hardware support. Ubuntu just keeps getting better.

Booting

When Canonical CEO Mark Shuttleworth introduced Ubuntu version 9.04, he highlighted the need for faster boot times. The noticeably improved boot performance under Ubuntu 9.04 is not because of any particular software. Instead, the developers used Bootchart (see Figure 2) to identify hangers and stoppers in the boot process and remove the culprits. The hardware detection phase was traditionally fairly slow, and Udev parallelization saves a couple of valuable seconds.



Phoronix, which does extensive performance tests on Linux distributions, compared Ubuntu 8.10 boot times with a development release, Ubuntu 9.04 Alpha 4. Testing on a netbook, Ubuntu 8.10 took 29 seconds to start, and the Jaunty Alpha 4 took just 21 seconds. For more details on the Ubuntu 9.04 boot performance, see the blueprint at the Launchpad site [2].

Cloud Computing

Many commercial cloud computing services are available today (for example, the well-known Amazon EC2 service), but third-party cloud vendors introduce risks. Issues such as security are important when considering a third-party vendor for such critical business needs.

This problem begs for an open source solution. Ubuntu 9.04 supports Eucalyptus [3], a set of tools designed to help you build your own cloud. Eucalyptus is interface-compatible with Amazon EC2, meaning anything you do on EC2 you can do with Eucalyptus.

Notifications

The most hotly discussed feature of Ubuntu 9.04 is notifications. Ubuntu's new notification system is similar to Growls in Mac OS and the Mumbles Projects for Gnome. A notification is not associated with an action, and you don't need to click anything to make it go away. A message appears on the desktop (Figure 3), and then, after the information is presented, it simply disappears. At least 35 applications support Ubuntu's notification system, including Network Manager, Pidgin, and Evolution. Popular applications such as Firefox still need to integrate notification support.

tu's notification system, including Network Manager, Pidgin, and Evolution. Popular applications such as Firefox still need to integrate notification support.

New Kernel

One of the most obvious changes with Ubuntu 9.04 is the new Linux kernel version, which includes numerous new drivers (for Ethernet and WLAN cards, filesystems, and webcams), and also many other updates.

At the top of the list for 2.6.28 kernel innovations is support for the new Ext4 filesystem, the heir to the previous Ext3. Ext4 is faster and can handle more files than Ext3. A file can be as big as the whole filesystem; partitions and volumes of up to 1 exabyte (equivalent of one million terabytes) are supported.

If you want to update your legacy Ext 2 and Ext 3 filesystems to Ext 4, you will need to run the *grub-install* command to let the grub boot manager correctly identify the kernel. Also note that you won't be able to mount Ext 4 partitions in environments with kernel versions prior to 2.6.28.

Ext4 is not without controversy. Testers have found is-




Figure 1: Members Ubuntu team gather in London to celebrate the Jaunty release.

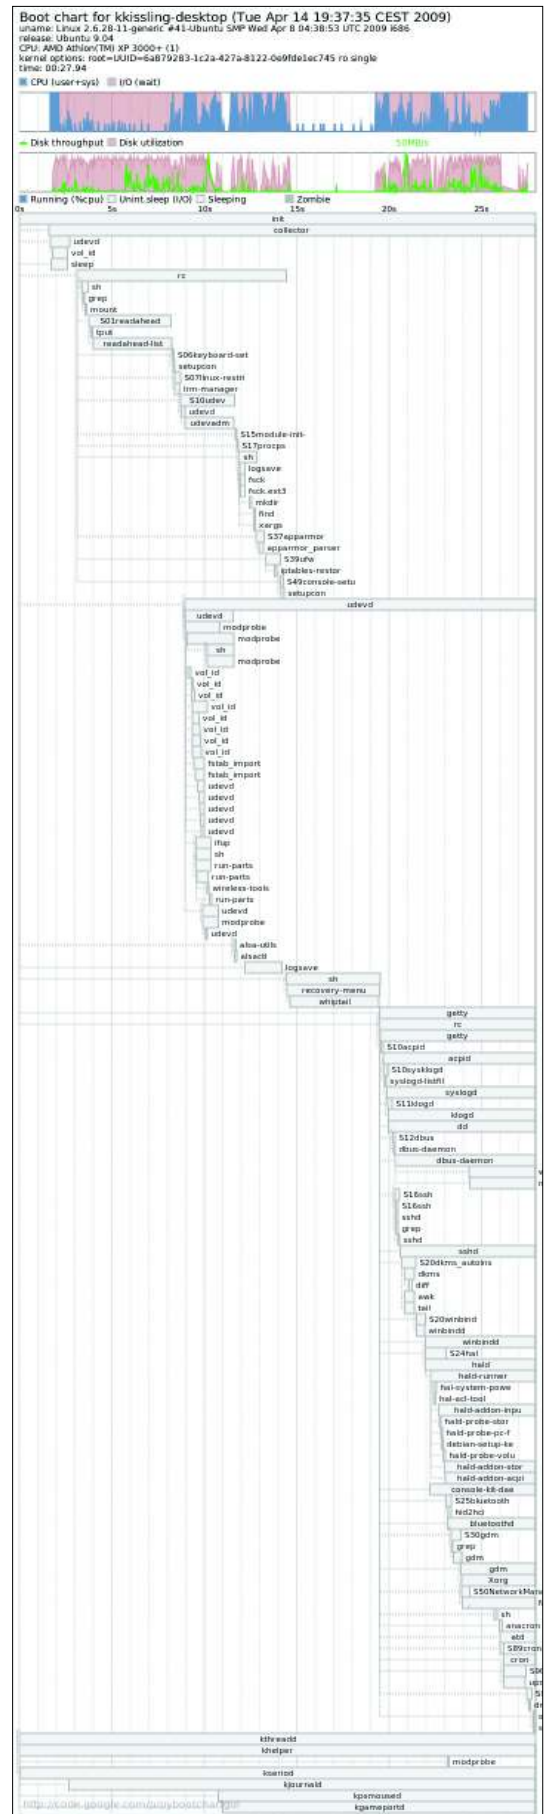


Figure 2: Bootchart lets Linux users measure the time required by active component during the boot process.

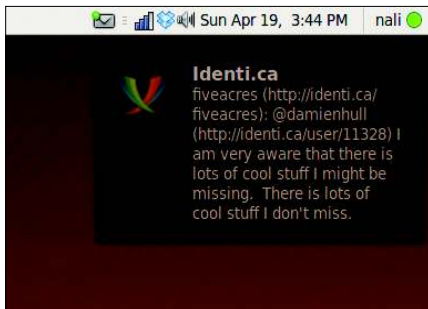


Figure 3: A notification from the microblogging service identi.ca.

sues with the way applications update files and how Ext4 writes to the disk. Discussions are ongoing in the community. In the meantime, Ubuntu 9.04 does not use Ext4 as the default filesystem for fresh installs, but it is an option if you want to partition a drive manually when installing from the Live CD.

The new audio drivers provided by the kernel now support “jack sensing” – that is, they auto-detect microphones and loudspeakers. Various code changes also improve Linux’s handling of DVB hardware, webcams, and various other devices.

Desktop

Many of the improved applications in Ubuntu 9.04 result from improvements to the Gnome 2.26 desktop, which offers simpler file sharing and better multiple-monitor support.

PulseAudio

PulseAudio brings a new sound server to Ubuntu that lets users individually set the volume for each sound application.



Figure 4: Thanks to PulseAudio, you can individually set the volume for applications.



Figure 5: The new 3.0 version of OpenOffice with PDF import and other features.

The software has a questionable stability record, and anybody interested in using it will need to install the *gnome-volume-control-pulse* package via the package manager. After restarting, you will see two small loudspeaker icons in the panel. You need to remove the one with the vertical slider. To do so, right click the icon, and select *Remove from panel*. Right clicking the remaining icon and selecting *Preferences* in the drop-down takes you to the PulseAudio configuration dialog. You can now set the volume for individual programs in the *Applications* tab (see Figure 4).

OpenOffice

Version 3 of OpenOffice (Figure 5) has been around since October 2008, but it was released too late to make the Ubuntu 8.10 freeze. The major Version 3 release included with Ubuntu 9.04 introduces a number of changes. For example, OpenOffice now supports PDF import, thus giving users (restricted) PDF editing ability. The developers point to the extensibility of the office suite via extensions that work in a similar way to Firefox add-ons. Users can download an exten-

sion to add its functionality to the office software. Thanks to improved support for OOOXML – Microsoft’s XML file format – the software now does a much better job of importing documents from Office 2007, but OpenOffice still fails to import highly complex OOOXML documents.

Janitor Tool

Ubuntu 9.04 also introduces a janitor tool. This helpful garbage collection service disposes of files and packages that you no longer need (Figure 6). You might need to install 10 independent libraries just to add a single application package. If you delete the application, the Synaptic package manager will still keep the other library packages on the hard drive. The Janitor can identify and resolve this kind of problem. But there has been some criticism of the tool – at least in the early alpha versions of Ubuntu 9.04. The janitor reportedly deleted some software that users had installed manually. A cautious approach is your best bet.

Ekiga

Ekiga 3.0 comes with a free video telephony tool for Linux. Although Ekiga is not compatible with the popular Skype, it is a useful replacement if the people you want to contact all have Linux. The developers have completely reworked the program’s interface, at the same time improving the video quality. Ekiga 3.0

collaborates with Evolution and will even import Evolution address books.

Pidgin

The Pidgin Instant Messenger tool, formerly Gaim, is a useful choice for occasional Internet chatting. Pidgin lets you talk to users with messengers such as AIM, Google Talk, ICQ, MSN, IRC, and Yahoo, although the Pidgin client will not necessarily support video functionality for all these providers.

Kubuntu and Xubuntu

KDE users who congregate on the Kubuntu side of the landscape will also find many changes with the latest release. Kubuntu 9.04 [4] comes with KDE 4.2.2. Amarok has always been one of the most popular apps for KDE, and many people consider it the best media player available. Kubuntu 9.04 includes Amarok 2, a complete overhaul of previous versions. The user interface lets you view lyrics and albums from the same artist together. Much of the work for Amarok 2 went into integrating online services such as Last.fm, Jamendo, Magnatune, and SHOUTcast.

Quassel is now the default IRC application in Kubuntu. IRC is a great way to get support and connect with the rest of the Ubuntu community. (Most Ubuntu

projects have their own IRC channels.) What makes Quassel interesting is that it can run in a distributed manner. The core component of Quassel can run continuously on a server and stay connected to various IRC networks. Client Quassel components can then connect to the core and use IRC normally. When the client disconnects and reattaches later, the IRC sessions will continue just as if the client had never disconnected.

Not to be left out, Xubuntu [5] also arrives with some promising new features. Ubuntu's Xubuntu variant is a favorite of users who want a streamlined desktop. Xubuntu 9.04 includes the latest Xfce 4.6 desktop environment, which offers better desktop integration. Thunar, the file manager, now has context menus that let you open applications, exit user sessions, or access help documentation. The mixer now works with GStreamer, which means that it supports multiple sound systems. Other features include a new session manager and a new settings manager.

Support and Community

If you're looking for advice online, the Ubuntu Forums [6] offer a first level of support for many Ubuntu newcomers. The forum system even includes a subforum for new users: Absolute Beginner Talk. Other forums range from gaming to networking and wireless. Threads are monitored continuously by volunteer moderators to make sure no one is posting malicious commands.

One aspect of Ubuntu that is different from other distributions is the presence of local communities (LoCos). Similar to LUGs (Linux user groups), LoCos advocate, promote, translate, develop, and otherwise try to improve Ubuntu. LoCos are not meant to replace LUGs but, instead, they work with existing LUGs to evangelize and support users. Almost 200 LoCos have formed all over the world. If you are looking for a way to connect with Ubuntu users in your area,

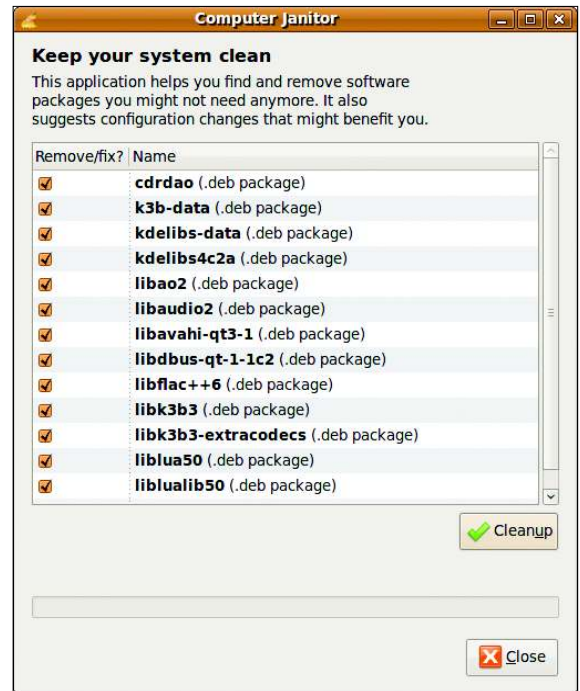


Figure 6: The new Ubuntu Janitor is designed to handle orphaned files, but its approach might be too aggressive at times.

check out the LoCo team list in the Ubuntu wiki [7].

Read On

Welcome to *Ubuntu User* magazine. Read on for news, views, and informative discussions of some of the best tools of the Ubuntu environment. You'll learn about virtualization, firewalls, Windows integration, and the steps for building your own website.

If you are new to Ubuntu, you might want to start with our exclusive Discovery Guide section, which shows you how to set up your Ubuntu system and helps you put Ubuntu to work on some common tasks. ■

Don't Zap

Long-standing Ubuntu users will probably notice the "Don't Zap" option. Whereas previous versions let users quit the desktop by pressing `Ctrl+Alt+Backspace` as a short-cut to the Gnome login window, this option is no longer available. Many users will probably consider this change a disadvantage. The new approach is to use the `Alt + PrtScr + K` keyboard shortcut, which the kernel has supported for many a moon, to handle the job. This shortcut also restarts sleeping desktops.

To reestablish the previous shortcut, you need to reinstate the original `/etc/X11/xorg.conf` configuration file status. To do so, press `Alt+F2` and type `gksu gedit` to launch an editor with administrative privileges. Open the file and add the following lines to `xorg.conf`:

```
Section "ServerFlags"
    Option "DontZap" "False"
EndSection
```

Then restart your system.

INFO

- [1] Ubuntu: <http://www.ubuntu.com/>
- [2] Boot time details: <https://blueprints.launchpad.net/ubuntu/+spec/jaunty-boot-performance>
- [3] Eucalyptus: <https://help.ubuntu.com/community/Eucalyptus>
- [4] Kubuntu: <http://www.kubuntu.org/>
- [5] Xubuntu: <http://www.xubuntu.org/jaunty/countdown>
- [6] Ubuntu Forums: <http://ubuntuforums.org/>
- [7] Ubuntu LoCo Teams: <https://wiki.ubuntu.com/LoCoTeams>



Meet Ubuntu's Community Manager

DIVERSE UBUNTU UNIVERSE

Ladies and gents, boys and girls, I would like to give you a very warm welcome to my brand spanking new column.

BY JONO BACON

today, it limited the scope of who could get involved.

Today the story is far different. Every day we hear wonderful tales and experiences across the free software world about the diverse range of people joining us. We hear of people exercising their passions, whether it be in documentation, translations, advocacy, design, testing and other areas we have not thought about. As we have become more accepting of this diverse range of contributions (and some beards), it has been fascinating learning how people work together to make different contributions welcome. This has become a really interesting focus of my work: understanding how people collaborate and optimize for the nuances in that collaboration.

I have always been a firm believer that different brains naturally gravitate toward different things. As an example, in the programming sense, I naturally gravitate (my quite awful programming abilities) toward Python. My friend Dan naturally gravitates towards Perl. While our abilities are similar, my brain simply doesn't work with Perl, and his won't work with Python for a variety of reasons, including expectations, experience, technical cultural differences, and workflow. These considerations are important when assessing how communities are formed.

Part of my work has been in understanding the subtle differences and connections between different groups of people and figuring out ways to develop the potential around our findings. When we identify these similarities and patterns between different disciplines, it often leads to interesting results. A good example of this is communication: Soft-

ware developers love mailing lists, but generally don't get on very well with forums. Although there is nothing specific that excludes developers from forums, many tend to prefer mailing lists because the content is delivered directly to them. Forums are the antithesis of this ethos and instead require you to hammer that refresh button until you see new content. Knowing these nuances helps us to build better communities because we don't force contributors into boxes that they don't naturally belong in.

So far, in the Ubuntu community we have tended to some of these roles pretty well. We have a thriving community engaged in packaging, testing, documentation, translations, and advocacy, but I am particularly interested to inspire new avenues of participation. In my most immediate radar I see design and user interface testing as particularly interesting targets. Design is a very visual art-form, and designers collaborate using visual gestures and imagery.

This is just one such example of the challenges that face us, and through the following series of columns I am looking forward to sharing my experiences and thoughts on how to inspire and enthuse this new generation of free software contributors.

Before I run, I just wanted to let you know that I am really keen to hear your feedback, so feel free to get in touch with me at jono@ubuntu.com. ■

When I was first offered the opportunity to write a column, I scratched my head and wondered what on earth I should cover. The problem was certainly not a lack of ideas, but quite the opposite: There is so much going on in the Ubuntu community that I wanted to use my column as a means to showcase many of the great stories that happen each and every day. I have the pleasure of being exposed to oodles of these stories, and I can guarantee you that this is no help when narrowing something down for my column. Fortunately, it seems the topic was staring me in the face all the time, and deeply connected to all these stories: Diversity.

The Future of Participation

I first got involved in Open Source 11 years ago, and back then it was a very different place. If you wanted to contribute to the community, you were predominantly welcomed if you were a programmer. Filling these shoes had high expectations: the typically bearded individual was assumed to know how to program in C and have a strong familiarity with a range of obtuse technical tools.

While these heady days produced some incredible code and laid down the foundations of the Ubuntu we know


INFO

[1] Jono Bacon's homepage: <http://www.jonobacon.org/>

[2] Art of Community Online: <http://www.artofcommunityonline.org/>

Exploring Ubuntu's virtualization tools

GET VIRTUAL



Marina Strizhak, 123rf.com

Why tie yourself to the hardware? We'll show you how to operate virtual environments from your Ubuntu desktop.

BY MARCEL GAGNÉ

After just loading your copy of the shiny new Linux release Ubuntu 9.04, things are going well. Now that you've read the other articles in this issue, you're starting to think you've found a pretty good home. But even when you love your home, once in a while you want a holiday – which means loading up another Linux distribution to play with.

Of course, it might not be all play. Sometimes you need to test another operating system, or you might need to run

Windows as well as Ubuntu. Wouldn't it be nice if you could run all those things on your current computer, without having to reboot or reinstall? You can – with virtualization.

Virtualization software creates a virtual computer on your computer, on which you can install any other operating system, whether it be Ubuntu, Red-Hat, Mandriva, or yes, even Windows. From the perspective of the guest operating system, the virtual machine is a real machine with real hard drives, a real

network card, real video hardware, and so on. When you boot that virtual machine, it will appear as though you are booting a real machine.

Virtualization Tools

Your Ubuntu system doesn't have any of this software installed by default, but getting it is easy. Furthermore, several virtualization packages are available. These include the command-line programs QEMU [1] and KVM [2], as well as one of my favorites, the powerful graphical VirtualBox [3]. To install these programs, select them via Synaptic or the *apt-get* command:

```
apt-get install qemu kvm  
virtualbox-ose
```

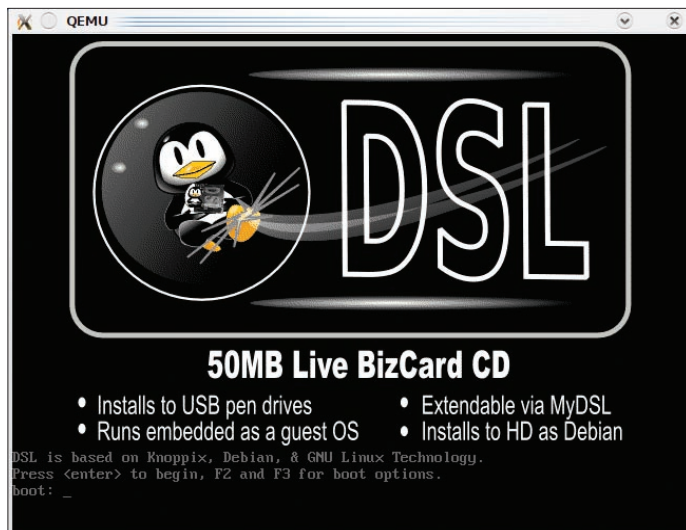


Figure 1: QEMU virtual hardware boot.

Of the three, I prefer VirtualBox, but let me show you how easy it is to run a virtual machine using QEMU or KVM. KVM, the Kernel Virtual Machine, is a replacement for QEMU. It works by taking advantage of processor-based hardware virtualization technology. On the Intel platform, this is known as VT; on AMD processors, it is called Pacifica. What it means to you is that you can take advantage of KVM to run your virtual machines with substantially better performance than you can with straight software virtualization. The catch, as you might expect, is that not every machine out there supports hardware virtualization. To find out whether your system is a candidate, execute this little snippet of code from a terminal window or shell:

```
grep -E ^flags.*(vmx|svm) /proc/cpuinfo
```

If it returns the string `vmx` or `svm`, your processor is ready. If not, you have to stick with the software-based virtualization of QEMU. The KVM commands are pretty much interchangeable with those of QEMU, as are the install and operating system images. In fact, KVM is based on QEMU.

Next, I'll take a look at how QEMU works. For

this first demonstration, I'll install a tiny distribution called Damn Small Linux (DSL) from an ISO image downloaded from the project website [4]. Because DSL is a cute little distro with minimal space requirements, I'm going to create a relatively small disk

image (a virtual hard disk) for it to live on with the `qemu-img` command:

```
qemu-img create dsl.img 256M
```

The system returns with the following reply:

```
Formatting 'dsl.img', fmt=raw, size=262144 KB
```

The preceding command creates a raw-formatted disk image by default. A few different image formats are available, most notably `qcow2`, which is a portable image format that is useful if you want

to install a non-Linux OS – like Windows XP, for example.

The next step is to install Linux into this disk image, which I do with the `qemu` command:

```
qemu -cdrom dsl.iso -hda dsl.img -m 256 -boot d
```

Several interesting things are happening here. For starters, the `-cdrom` parameter is the path to the CD-ROM image from which you are installing your distribution. If you were using a physical CD-ROM, that path would likely be something like `/dev/cdrom`. The next parameter, `-hda`, defines the path to the disk image I just created; then, it's followed by the `-m` switch, which, in this case, allocates 256MB of RAM to the running session. Finally, `-boot` identifies the boot drive (the CD image), which is the so-called D drive. As soon as I press Enter, the virtual hardware boots (Figure 1).

The disk you created and the CD-ROM are represented as physical devices at boot time, as is the default video card and the BIOS. In fact, you can even interrupt the boot process and bring up a device boot menu. If you press Enter at this point, the system will go through its normal boot process, which, in the case of a Live CD, usually stops to let you decide what to do as far as booting or installing. DSL is built on Knoppix, so

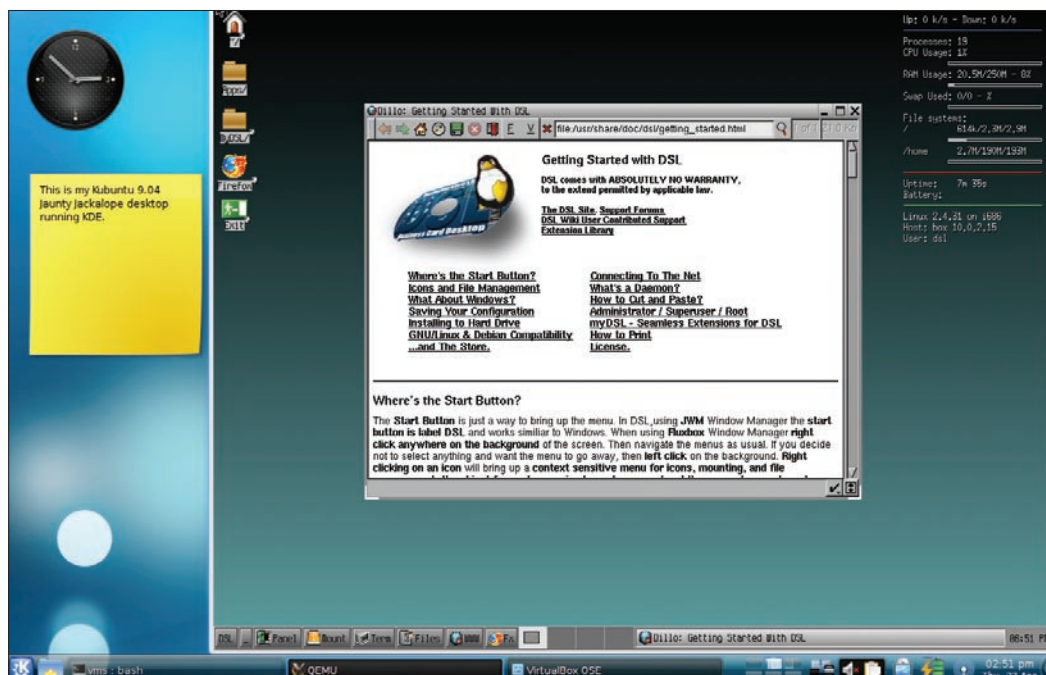


Figure 2: A Kubuntu Jaunty Jackalope desktop with an instance of Damn Small Linux booted and running on top.

users of the venerable Live CD will find some of this familiar. For everyone else, just press Enter.

Everything that happens from this step forward should happen as it would if you were running any other Linux distribution from a Live CD. The guest operating system will boot, or install, normally. Depending on the speed of your machine and the amount of memory you have at your disposal, the process could be peppy or quite slow. Remember that this virtual machine uses whatever portion of your host system's resources you allotted. That means less for your host system, both in terms of memory and processor resources. Also, if your processor allows KVM instead of QEMU, the performance will be substantially higher.

A few seconds (or minutes) later, your virtual machine is up and running (Figure 2).

KQEMU

Before, I suggested that running QEMU instead of KVM means you aren't privy to processor-level acceleration, which means you are emulating everything in memory without the benefit of kernel acceleration. When you are running a little tiny distribution like DSL, that's fine, but if you want to try a different Ubuntu, openSUSE, Fedora, or even Windows, you will feel the pain pretty quickly. In this case, the KQEMU hardware accelera-

tion module comes into play. The catch is that it's not installed by default, and you can only install the source for it. Yes, that means you have to build it. Luckily, that's not complicated:

```
sudo apt-get install module-assistant
module-assistant
sudo module-assistant prepare
sudo module-assistant -f get qemu-source
```

The result is a kernel module that dramatically speeds up your QEMU sessions, giving you performance on par with KVM. Now, even with QEMU, you can run and install more demanding operating systems, like Ubuntu. For the pure fun of it, I loaded up Ubuntu Jaunty Jackalope (with the Gnome desktop) and ran a full implementation of it on my Kubuntu desktop (Figure 3).

Before I move on, I'll share a few other parameter with you. For starters, you might add *-no-acpi* to disable this in the virtual machine. Also, you can use the *-localtime* parameter to tell QEMU to run your session with the local machine time as opposed to the default UTC. Remember the kernel acceleration? Try *-kernel-kqemu*. In addition are network parameters, USB parameters, file-sharing parameters (SMB), devices parameters, and more. Because you are only seeing




Figure 4: Your first step is to tell VirtualBox what guest operating system will run on this VM.



Figure 5: Unless you defined a previous virtual disk, you'll need to create one. Because it is possible to reuse disk images, previous images will be listed here.

the basics, make sure you check out the man page for the *qemu* command. Also

remember that my settings for this example imply a distro with a very small footprint.

VirtualBox

Another great piece of virtualization software is VirtualBox, an open source package freely distributed under the GPL and distributed by Sun Microsystems (who, as I write this, has been purchased by Oracle Corporation). VirtualBox is a program I use every day and one I recommend highly. With VirtualBox, I can load an operating system that you might still want access to: Windows XP. Of course, you do need a licensed copy of

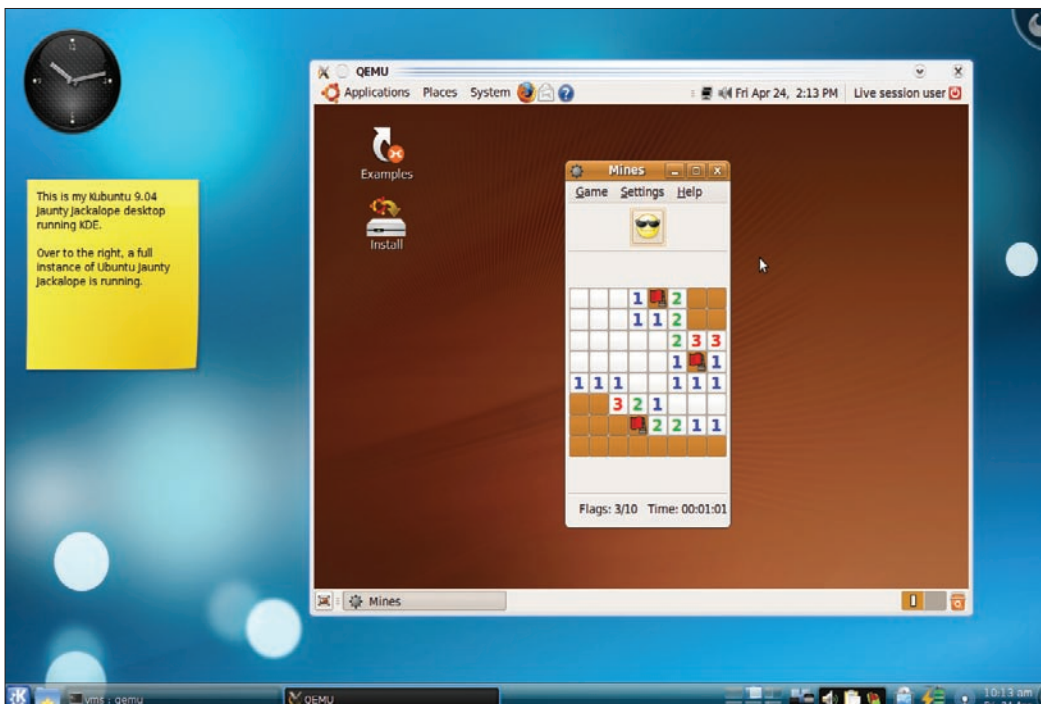


Figure 3: The best of both worlds: Ubuntu's Jaunty Jackalope 9.04 running in QEMU under Kubuntu.

Microsoft Windows to proceed.

When you start VirtualBox for the first time – assuming you already have Virtualbox installed via either Synaptic or apt-get – no machines will be running in it. Think of it as a blank slate or, better yet, a new computer with a blank hard drive waiting for your favorite distribution, or, in this case, Windows XP.

Before you can install Windows, you need to create your virtual hardware. To do this, click *New* on the VirtualBox toolbar. A nice welcoming message explains that the wizard will guide you through the various steps. When you click *Next*, you will be asked to decide what kind of virtual machine you plan to create (Figure 4).

To start, give it a name. For example, I've chosen to call mine the very unoriginal *WinXP*, but you can call it whatever you like. Next, select the operating system (Microsoft Windows) from the drop-down list, then identify the version (Windows XP). The list of operating systems covers a lot of ground, including Windows, Linux, Solaris, BSD, and even IBM's old OS/2, and each of these operating systems have many versions. The Windows support goes right up to the new, not-yet-released Windows 7. When you click *Next*, you will be asked to select the amount of RAM you want to dedicate to the VM.

Resources

Deciding how much memory to allocate is a bit of a balancing act. On one hand, most modern operating systems have minimum requirements. On the other hand, you have to take into consideration how much memory your Ubuntu host system has. Either enter the amount into the text box or use the slider.

Next, your new machine needs a hard drive on which to load XP, so you must create a virtual hard disk (Figure 5). Because this is your first time out, the primary master (the main disk) will be

listed as *<no media>*. Clicking *New* creates a disk image.

Clicking *Next* in the new welcome message for the virtual disk wizard leads you to the dialog from which you can choose one of two disk image types: a fixed-size storage medium or a dynamically expanding disk image. The footprint of a dynamically expanding image is minimal to start and expands as needed. In comparison, the fixed image takes up whatever space you give it right from the beginning and runs up against a wall when that space is used. Unless you have good reason to do otherwise, select *Dynamically expanding storage* and click *Next*.

Having made your decision on storage options, you can now define the size and

location. The location, by default, is your machine name (my WinXP in Figure 4). In effect, this is a folder on your disk in which the machine resides. Next, you allocate the space either by typing it into the text window or with the slider. A size for this disk will be suggested according to the operating system you are loading. On the basis of your expected needs, accept the default or choose a size. After you click *Next*, you can review your selections. If you are happy with the choices you've made, click *Finish*.

Now you are back to the virtual hard disk selection screen you saw in Figure 5, but this time, your new disk image is selected for you. All you have to do is click *Next* to continue. This opportunity is the last you'll have to review everything you've done to date. VirtualBox will remind you of your chosen machine name, the memory allocated, the OS type, and the type and size of disk you've chosen. When you click *Finish*, your virtual machine is ready. In the main VirtualBox window (Figure 6), your new WinXP machine appears in the list of machines (on the left), ready to be loaded.

Remember, this is a clean computer, with nothing installed. To the right of the machine list are the parameters associated with the selected machine. Each of the blue labels can be clicked to config-




Figure 6: Your new machine appears, ready to be started.

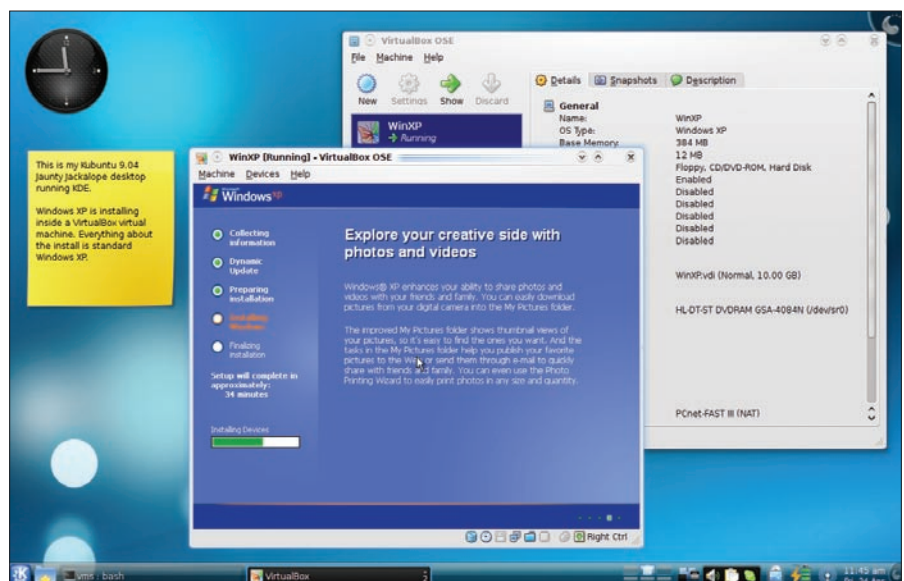


Figure 7: The new virtual machine is live, and Windows XP is installing normally. With VirtualBox running on my Ubuntu KDE desktop, the XP installation is in the foreground.

ure the resource listed. Notice that you can even go back and change the virtual machine specifications, such as memory. Other labels let you define what physical resources (on your host machine) the virtual machine has access to. For instance, to give your virtual machine access to the physical DVD-ROM drive in which the Windows XP install media is located, click on the *CD/DVD-ROM* label.

In the resulting dialog box, check the *Mount CD/DVD Drive* box. If you have more than one drive, select the one you want from the drop-down list. The radio button labeled *ISO Image File* is particularly interesting if you are loading another Linux distribution because you can install it directly from the hard drive, skipping the need for a physical disc (as is the case with Windows because free Live DVDs don't exist). When you click *OK*, you'll find yourself back at the main VirtualBox window with the *CD/DVD-ROM* label now indicating the host computer's physical drive. Now you are ready to install Windows XP.

Installing the Guest

First, make sure the virtual machine is selected, then click the *Start* button. VirtualBox will let you know that the auto-capture keyboard feature has been activated. This includes your mouse as well. What this means is that the virtual machine has full access to your mouse and keyboard. To "uncapture" your mouse and keyboard, press the right Ctrl key on your keyboard. To acknowledge the message, click *OK*, and your new computer

will boot from the Windows install disc. What follows is a standard Windows XP installation (Figure 7): You accept the license agreement, your virtual disk is formatted, and Windows does its thing.

Somewhere in the process, you'll need to enter your license code, answer some questions, and so on. Eventually, the installation will complete and your new Windows machine will be up and running. After you attend to standard Windows stuff, such as adjusting screen size and making other modifications to your running machine, you'll have a full Windows XP implementation running on top of your Ubuntu (or Kubuntu) desktop, with access to everything you normally use under Windows (Figure 8). Instead of emulating Windows, you are running a real, albeit virtual, PC loaded with Windows XP.

Although this might sound like a happy ending, in which the two seemingly irreconcilable operating systems are working together on the same hardware, VirtualBox has Guest Additions that can improve the relationship. Once you've had to hit the control key a few times to recapture your keyboard and mouse or you find that the default screen sizes provided by Windows just don't work well on your widescreen notebook, you'll be happy to learn that even these issues can be resolved. To install the Guest Additions, release your keyboard and mouse (right Ctrl key), then click the *Devices* menu on the virtual machine and select *Install Guest Additions* from the menu (Figure 9).




Figure 8: Windows XP running on my Kubuntu desktop - fully integrated and ready for action.

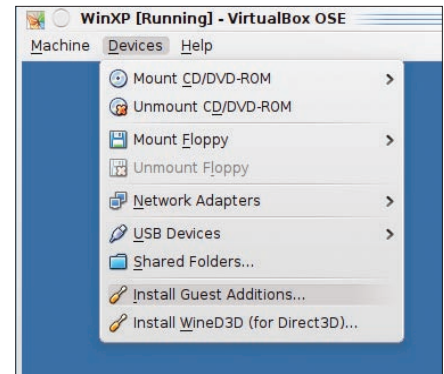


Figure 9: For nearly seamless integration into your Ubuntu desktop, install VirtualBox's Guest Additions into your running XP machine.

After you've followed the instructions for installing the guest additions and restarted your machine, your experience will be vastly different. Just by clicking on a window inside Windows, you can change the focus from your Ubuntu applications to your Windows desktop. Your mouse will be able to sail across your desktop from Windows to Ubuntu and back again. Furthermore, you will be able to resize Windows to fit whatever geometry makes sense to you and your Ubuntu desktop. Finally, you will be happy to know that these guest additions are not specifically a Windows thing and are available for other operating systems as well.

Conclusion

Happily ever after is where this story ends. By running Ubuntu, you have a superior operating system that is faster, more reliable, and more secure than Windows. But you can have your Ubuntu and your Windows too. Best of all, you aren't limited to just Windows. For example, you can try other Linux distributions, and even BSD or Solaris, from the comfort of your existing Ubuntu installation without reinstalling.

Yes, you can have it all - with a little virtualization. ■

INFO

- [1] QEMU: <http://www.nongnu.org/qemu/>
- [2] KVM: http://www.linux-kvm.org/page/Main_Page
- [3] VirtualBox: <http://www.virtualbox.org/>
- [4] DSL: <http://damnsmalllinux.org/>

Gaming on Ubuntu

PLAYGROUND



Vasileios Akroinoqou, 123RF

Gaming on Ubuntu is a reality. In this article we introduce 10 games that you can install without breaking a sweat. **BY TIM SCHÜRMANN**

Installation

All of the games I will be looking at here are easily installed via *Applications | Add/Remove*. When you get there, set *Show to All available applications*, then browse the *Games* section. If you can't find what you're looking for, the *Search* box might help.

In contrast to Windows, Ubuntu is not generally regarded as an El Dorado for computer gamers. But this widespread opinion is actually unjustified, as the Synaptic package manager goes to show with its choice of more than 1,000 packages in the *Games* section. Many games promise hours of fun when your daily chores are done, and it's totally free of charge. On the following pages, I will present a short selection of very popular Linux games. But beware, some of them are addictive.

1 Freeciv

The Freeciv strategy game is a classic Linux game. Just like its role model, Civilization, the player guides a civilization from ancient to modern times en route to prosperity. This endeavor is not just a question of exploring new territory, negotiating with other tribes, and conducting research, but also of surviving the odd war or two. The game is related to Civilization II, but in many cases, you can customize it to suit your needs.

Freeciv: <http://www.freeciv.org>

Alternative: *FreeCol*

2 Lincity and LinCity-NG

If you have always fancied a term in office as the mayor of a city, Lincity is just what the doctor ordered. Another Linux original, the game is modeled on the popular SimCity. As an urban planner, you are given the task of creating a prosperous city from a piece of vacant ground. For this to happen, you have to assign land for construction, build roads, and collect taxes. Each decision you make has an effect on urban life: Houses in the suburbs will attract wealthier citizens, whereas ghettos will develop in the vicinity of heavy industry, requiring expensive policing. An addition to the fairly frumpy-looking Lincity is a modern version called LinCity-NG.

Lincity variants: <http://lincity.sf.net> and <http://lincity-ng.berlios.de>

Alternative: *OpenCity*

3 Beneath a Steel Sky

If you like solving puzzles and are into science fiction, you will not want to miss Beneath a Steel Sky. This old-school adventure lets you direct the protagonist, Robert Foster, through Union City, a city built on iron and steel. Supported by a little robot, Joey, you battle the merciless, totalitarian state and in the process discover a dark spot in your own past. Beneath a Steel Sky started life in 1993 as a commercial game for MS-DOS, which explains the slightly pixelated, but atmospheric, graphics. On Ubuntu, it requires the ScummVM utility, which has given many other classic adventures a new lease on life on Linux.

ScummVM: <http://www.scummvm.org/downloads.php>


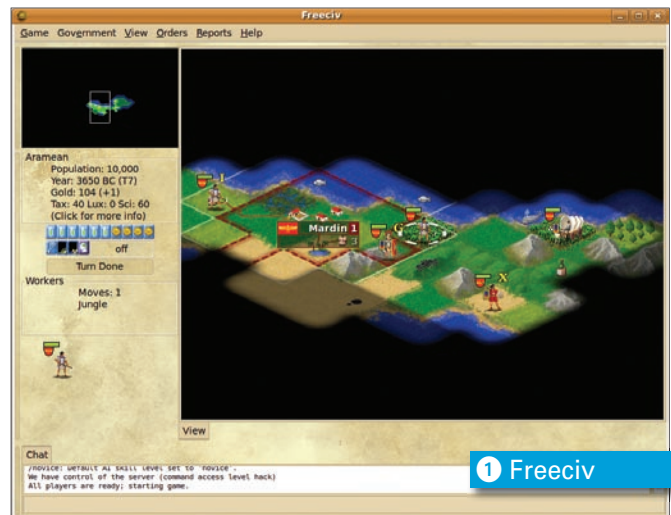
Alternative: *Flight of the Amazon Queen*

4 FlightGear

The makers of FlightGear were looking to create a realistic and extensible flight simulator. The graphics can't compete with the commercial counterparts, but the flight characteristics and environmental conditions are all the more realistic to make up for this. The sheer number of options and settings are likely to put off occasional flyers, though. FlightGear does not just simulate every single on-board instrument, but every single lever on board your choice of plane. Before you even dream of taking off, plan a couple of hours practice. A giant, world-wide community continually turns up with new maps and new flying material.

FlightGear: <http://www.flightgear.org>

Alternative: *ACM Aerial Combat Simulator*



5 Battle for Wesnoth 2



5 Battle for Wesnoth

Battle for Wesnoth is another turn-based strategy game modeled closely on the erstwhile favorites Master of Monster and Fantasy General. In a fantasy world divided into hexagonal cells, players move their characters, each with their strengths and weaknesses, and gain experience in the course of battle in a fashion similar to role-playing games. Type of attack and weapon, positioning, time of day, and terrain are all important to game play. A good introduction helps you take the first steps, and thanks to the purposely simple rules, newcomers can come to grips with the game in next to no time. The graphics never cease to amaze with excellently rendered manga-style images. Playing against human competitors is particularly enjoyable, and the map editor included with the game lets you define your own campaigns, which you can swap with other fans on the Internet.

Battle for Wesnoth: <http://www.wesnoth.org>
Alternatives: *Advanced Strategic Command, Glest*

6 TORCS




6 TORCS

Motor sports fans can hit the virtual throttle in The Open Racing Car Simulator, or TORCS for short. At the wheel of a modern race car, such as a McLaren F1 or a Mercedes CLK DTM, you can race over intelligently laid-out courses. Among the more than 20 tracks to choose from are high-speed courses on smooth tarmac, dusty rally courses, and street racing. TORCS simulates the response of more than 50 different vehicles based on a number of parameters, such as suspension, aerodynamics, and driver error. Attacks by opponents are reflected by a simple damage model. Because the game does not support multiple players, you will have to make do with racing the computer.

TORCS: <http://torcs.sourceforge.net>
Alternatives: *SuperTuxKart, Trigger*

7 FooBillard



7 FooBillard

FooBillard is another sporting challenge. This pool/billiards/snooker simulator offers attractive 3D graphics, realistic physics, and a simple computerized opponent. As an alternative, you can play human competitors over a network. A selection of games is available, including 8 ball, 9 ball, carambole, pool, and snooker. And if you happen to own a pair of 3D glasses, you can switch to anaglyph mode and experience a fascinatingly real 3D view of the table. Table textures, ball highlights, and shadows add to the realism of the game.

FooBillard project: <http://foobillard.sunsite.dk>
Alternative: *BillardGL*

8 SuperTux

If you are looking for an existing jump and run side-scrolling game for the kids, SuperTux is perfect. Instead of an overweight plumber, as in Nintendo's Super Mario series, you help a small penguin negotiate pitfalls and cuddly monsters. The latter typically can be knocked out by jumping on their heads. Bonus articles give the hero super powers. The game currently has more than 100 levels, to which you can add your own creations, thanks to the level editor.

SuperTux: <http://supertux.lethargik.org>
Alternative: *Secret Maryo Chronicles*

9 Vega Strike

Vega Strike takes you into the realm of outer space. This mix of action game and simulation sees the player set off as the owner of a small spaceship, very much in need of a refit. Of course, a refit is going to cost you money, and to earn it, you will need to trade with foreign planets, fight with or against factions, or indulge in a spot of piracy. Just as in its role models *Elite* and *Freelancer*, Vega Strike gives you an open environment. The universe, comprising more than 5,000 solar systems, is populated by 39 factions, all of which are tied up in alliances – if you are unlucky, you end up in the middle of a war. A kind of “warp” engine, called a SPEC drive, allows you to travel quickly throughout the solar system. Jump points allow you to travel between star systems – if you have a jump drive, that is. The game generates most of the planets randomly, and it's a completely new game on each restart. Many weeks of gaming pleasure are guaranteed en route to fame and power.

Vega Strike: <http://vegastrike.sourceforge.net>
Alternative: *Oolite*

10 Nexuiz

Ubuntu has no end of classy but free multi-user shooters. Nexuiz, a three-dimensional deathmatch game, is just one of them. The idea of the game is to take out the enemy in a team or on your own, to capture flags, or to secure waypoints. Multiple game modes are available. Just like the *Unreal* series, Nexuiz is designed for fast shoot-outs on the network; single-player mode, with more than 20 levels, is just for learning the ropes. The emphasis is on speed and intense action. Numerous, intelligently designed and attractively illuminated levels promise variety and strategically challenging duels.

Nexuiz: <http://www.alienstrap.org/nexuiz>
Alternatives: *OpenArena, Sauerbraten, Alien Arena, Tremulous, Warsow*



8 SuperTux



9 Vega Strike



10 Nexuiz

An interview with the Ubuntu Kernel Manager

JUGGLING JACKALOPES

One day before the official release of Ubuntu's "Jaunty Jackalope," Ubuntu User magazine interviewed Ubuntu's Kernel Manager, Pete Graner. Pete explains how a kernel grows into a Jaunty Jackalope and looks toward the next release of Ubuntu. **BY JOE CASAD**



Figure 1: Pete Graner (Ubuntu Kernel Manager), Mark Shuttleworth (Ubuntu Founder), and Jono Bacon (Ubuntu Community Manager) celebrating at the April 23rd Jaunty Release Party in London.

What does a busy, globe-trotting Ubuntu Kernel Manager do one day before Jaunty Jackalope hops out onto the world stage as the latest Ubuntu release? He tells *Ubuntu User* how the kernel team works, how the community contributes, and what the future looks like for the user-friendly Ubuntu.

Q How long have you been with the Ubuntu team, and how did you get involved with Ubuntu? What is your background?

A I've been with Ubuntu just about a year. This was my second Ubuntu release. Prior to that, I worked at Red Hat in various capacities, which include on-site SELinux Engineer at the National Security Agency (NSA) – I managed the Level 3 Support team – and my last position was the Base Operating System Engineering Manager for Red Hat Enterprise Linux, and by extension, Fedora. I started using and developing in Linux

while I was in the US Army as part of a Rapid Prototyping Project in about 1994. At the time, everything was Sun OS-based and we needed a lightweight, portable computer that could be used in flight, then jumped via parachute into the target area. The contractor selected Linux (Slackware) and i386 notebooks.

Q I get email from Ubuntu developers all over the world. How does development work? And I know a lot of volunteers are excited about Ubuntu. When it comes to kernel development, do you work directly with beta testers in the community?

A For the most part, we are one virtual team. The team is made up of two halves: Canonical engineers and community volunteers.

About 95 percent of the Canonical engineers are remote employees – that is, we all work from home. Canonical has a few offices – London, Boston, and Taipei – but for the most part, we don't work in

the offices. We work via Public IRC channels and mailing lists, and of course phone calls when needed.

The kernel team is broken down into 14 members, including myself. We are broken down into three focus areas. Stable Maintenance is responsible for our fielded supported releases, which include doing security updates and bug fixes. New Distro Development is stabilizing the kernel for the next development release, identifying and patching regressions from the previous releases, and distribution feature integration. Hardware Enablement is where we work with our partners, like Dell and HP, to make sure Ubuntu works with selected hardware models.

Ubuntu has a large community and during a development cycle we have public releases, so we don't really "Beta" test in the traditional sense. We begin testing at our very first development release – called Alpha 1 – through our Beta release and into the Release Candidate (RC).

We are always testing. The time between the releases is two weeks, so our users/testers are always getting new software through a rolling stream of fixes.

Q So Ubuntu is Linux. People are generally aware that you get your code from the venerable Linux community project known as Debian. Does the kernel come from Debian? Or do you work directly with the Linux kernel project?

A The kernel does not come from Debian due to the need to select a kernel version that will provide Ubuntu with the features it needs.

We have many cases where we will need a newer kernel due to our hardware enablement goals or particular fea-

tures that can't easily be back-ported to an older kernel.

At the Ubuntu Developer Summit (UDS), we make the decision on what kernel version we will use based on the feedback we receive from the other development teams. So, for example, at the Jaunty UDS we decided on 2.6.28. At that point, we import the kernel from the upstream Linux kernel and we begin to track the development there. This allows us to provide the current upstream development testing and, in the best case, bug fixes.

Q Do you start with specific goals for each release? If so, where do the goals come from? Do you all sit down together and talk about a vision?

A Before each release we hold a meeting where we get all the Canonical engineers, community developers, and interested parties together and discuss what the needs are. This is called UDS (Ubuntu Developers Summit). The sessions at UDS are normally at the individual development team level, except when you have things that cross subsystems in Ubuntu. The goals and features with working plans come out of this meeting. This year's UDS is being held in Barcelona, Spain, during the week of May 25-29.

Q How does the process work? Can you walk us through the time line of an average release cycle?

A We develop to a 6-month-time-based release, with a release every April and October. Once we start active development, we have Alpha, Beta, and RC releases every two weeks. So it would look like: (Alpha) Release Candidate 1 (RC1), UDS, then RC2, RC3, RC4, RC5, RC6, Beta, RC, Final Release.

This of course means we have numerous milestones we have to meet in the cycle. You can view the 9.10 Karmic Koala schedule online (<https://wiki.ubuntu.com/KarmicReleaseSchedule>).

Q Ubuntu has a reputation for focusing on the end user – making Linux seamless and accessible to the beginner. How do you solicit feedback from users? And how do you integrate that feedback into the kernel?

A Most end users should never even know the kernel is there. We try to pick the right kernel configs and driver defaults so that the user's experience is a good one.

We receive feedback all the time about what hardware does or does not work and where we can find drivers if they exist for new hardware, so I would say feedback isn't an issue. We try to provide the best hardware support possible. The feedback is discussed on our public mailing list and, depending on what it is, we may implement it in the next kernel update or it might get deferred to UDS, where we need a broader discussion.

Q I know there is a vision to bring Ubuntu into the server and corporate environment. Can you talk a little about recent work on that path?

A I would have to say Ubuntu Server has moved out of the vision phase. Also, I'd like to caveat this by saying I'm not tied deeply into the server team plans. With that said, Ubuntu Server is a very solid product, deployed in numerous environments today.

As far as recent work, much of the focus has been on making the default server applications come out of the box with sane defaults for the average use case.

Q What's up ahead? Are you already working on the next release? Can you give us a hint about what the emphasis will be?

A We've been working hard on reducing the amount of non-upstream code that we ship and have been investing time and energy into making upstream vanilla kernel builds available for interested users. We have opened the 2.6.30 kernel tree and are doing the usual patch scrub and trying to eliminate as many out-of-tree patches as possible. And smoke testing. We won't really start development in earnest until after the kernel version is selected at UDS. Until then, it's all prep work and things that are kernel-version independent.

For 9.10, we have a large number of items on the plate. We will be looking at the 2.6.30 or .31 kernel. Some of the items for discussion at UDS will be moving off of grub to grub2 for boot loading, improving power management, improving Suspend/Resume/Hibernate, encrypted swapfiles, improving WiFi and 3G support, and much, much more. ■

INFO

- [1] Ubuntu development: <https://wiki.ubuntu.com/UbuntuDevelopment>
- [2] Ubuntu testing: <https://wiki.ubuntu.com/Testing>



Configuring accelerated graphics cards in Linux

CARD TRICKS

Some graphics card vendors offer limited support for open source systems like Ubuntu. If you're looking for high-end effects, you might need to get involved. **BY KLAUS KNOPPER**



open source business models are endless, so I'll just say I'm strongly biased in favor of open technical specifications and open source, even as I write this article about installing proprietary drivers for graphics cards through the use of general Linux techniques.

If you're troubleshooting a graphics card in Ubuntu, you might also be interested in this month's "Answer Man" column (pg. 46), which includes an answer on installing an nVidia graphics card.

Graphics Under GNU/Linux

Most GNU/Linux distributions for desktop and notebook computers use the Xorg graphics subsystem, which is an open source implementation of the X11 API used by almost all programs that have graphical elements.

Xorg handles communications with graphical programs, and it also interfaces with hardware drivers to make access to one or more graphics cards stable in a multiuser/multitasking environment, providing a wide range of configuration options: transparency, hardware-based shading and rendering,

Igor Kofronov, 1231f.com

tion for writing drivers – to any third parties except those bound by non-disclosure agreements. This secrecy makes it very difficult to write open source drivers for such hardware. These vendors argue that this policy ensures a market advantage; yet, other vendors who provide open specifications and good support for Linux compete well, often gaining a reputation for well-supported hardware that is easy to use and widely deployed on a range of electronic devices. Philosophical and political discussions about proprietary vs.

Linux probably has the best overall hardware support because of its wide application on a variety of devices. However, some hardware vendors make a strategic choice to avoid revealing their specifications – including example code and technical documenta-




Figure 1: Glxgears measures the drawing speed of the graphics system.

support for multiple monitors, and direct communication between the program and the hardware using a kernel-provided direct hardware access feature.

This direct and unfiltered access to hardware through Xorg explains why careful programming of the hardware-dependent components, especially the driver modules for various chipsets, is crucial for stability of the entire system. Locking up the graphics card can result in an unusable system or even a total freeze of the operating system. Open Source chipset drivers that come with Xorg are therefore carefully reviewed for potential problems by developers, and every now and then, updates are released that provide new features or improve performance.

Acceleration

Many modern graphics cards provide on-board processing capabilities for graphics-intensive applications such as 3D computer games. Instead of letting the CPU perform the calculations and transfer pixels between memory locations, the graphics card does this independent of the main system with the use of special command sets. This feature is often called “accelerated graphics.”

Each chipset has its own command set and access method for acceleration. For cards for which vendors provide sufficient information, accelerated open source drivers are available at Xorg, such as some of the older ATI Radeon cards [1], as well as many Intel-based chipsets (especially those that are present in many notebooks and netbooks).

The degree of available acceleration varies because not every feature of every

chipset is fully documented. But a good hint for Compiz Fusion, the well-known 3D window manager that features transparent windows, magnification, and fancy desktop effects, is the output of the *glxinfo* utility (Listing 1), which you can find by installing the *mesa-utils* package.

In Listing 1, the line with *Direct Rendering: Yes* is the most significant hint that the card is supported in accelerated mode.

Some *GLX_ARB...* extensions listed in the full output are needed by Compiz, so direct rendering (DRI) [2] is not sufficient in all cases. However, with DRI, chances are good that games that use OpenGL/Mesa will run fine on this card (e.g., *tuxracer*, *neverball*, and *egoboo*).

For measuring the drawing speed of your system, *glxgears* is a nice tool that’s also present in *mesa-utils* (Figure 1). Once you start *glxgears*, it shows an animation sequence of gears while it tries to refresh the animation and count the frame rate reached in each five-second interval.

An ATI Radeon 9800 card on a Pentium II running at a 16-bit 1280x1024

resolution reaches about 1,100 frames per second in *glxgears* with the open source *radeon* Xorg driver module [3], whereas on an EeePC 701 (900MHz, Intel 915GM graphics chipset), the *intel* Xorg driver reaches at least 500 frames per second.

Composite and 3D Support

Composite is the “transparency” extension that lets the graphics system overlap multiple layers, creating an “invisible” appearance (Figure 2).

Both the ATI Radeon 9800 and Intel cards I tested ran fine with Compiz, once the necessary AIGLX and Composite extensions were configured in the Xorg configuration file */etc/X11/xorg.conf*.

Listing 2 shows some of the most common options for different graphics cards needed to get Compiz (and other programs that use composite effects) to work properly. Not all of them are necessary for each and every card; some of these options are now even the defaults on newer Xorg versions and don’t have to be present anymore (although their presence usually doesn’t hurt).

After adding these changes, you have to restart the X server, which will terminate any currently running local graphical desktop sessions and most likely show the login screen again. Before killing your X server with Ctrl + Alt + Backspace (without knowing whether your new configuration will work), you can test the changes by starting a second server on a different console.

As root, type

```
X :1
```

in a terminal window, which will start a new Xorg session on the next free console – you should see a gray grid and the X-shaped mouse cursor. If that works,

Listing 1: *glxinfo* Output

```
01 name of display: :0.0
02 display: :0 screen: 0
03 direct rendering: Yes
04 server glx vendor string: SGI
05 server glx version string: 1.2
06 server glx extensions:
07   GLX_ARB_multisample, GLX_EXT_import_context, GLX_EXT_texture_from_pixmap,
08   GLX_EXT_visual_info, GLX_EXT_visual_rating, GLX_MESA_copy_sub_buffer,
09   GLX_OML_swap_method, GLX_SGI_make_current_read, GLX_SGI_swap_control,
10   GLX_SGIS_multisample, GLX_SGIX_fbconfig, GLX_SGIX_visual_select_group
```

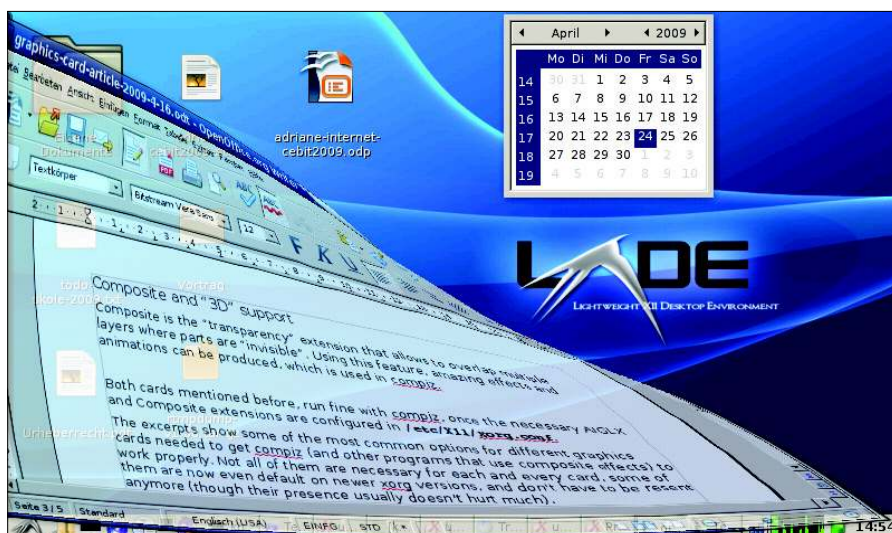


Figure 2: The transparency extension produces some interesting effects.

chances are good that your configuration is OK. With `Ctrl + Alt + Backspace`, you can kill the new X server and do the same (for restarting the X login process) on your first X session; better yet, logout and kill X in the login screen.

A more Unix-like way to achieve the same effect is to go to runlevel 2 or 3 by typing

```
init 2
```

(which is a text console-only runlevel in most distros) and restart the graphical session with:

```
init 5
```

Also, you need to do this when configuring a different chipset driver.

Choosing a Different Chipset Driver

If `glxinfo` reports *Direct Rendering: No*, the wrong chipset driver might be configured, so check `/etc/X11/xorg.conf` in the *Driver* subsection:

```
Section "Device"
...
Driver "fbdev"
EndSection
```

The framebuffer driver for Xorg, `fbdev`, is the one that always works if everything else fails, as long as framebuffer mode is enabled in your Linux kernel and possibly activated at startup.

To enable framebuffer mode, add option `vga = 891` to the kernel boot options,

which will give you a 1024x768, 16-bit framebuffer resolution that will be used by the `fbdev` Xorg driver. On startup, you will see the Linux penguin logo with a glowing border if `highcolor` framebuffer mode is on. Although this mode works well with most graphics cards, it is never accelerated, and the previously mentioned options for `Compiz` just won't work.

Try changing the *Driver* option in the *Device* section of `xorg.conf` to the best matching chipset driver, like `ati`, `intel`, or `mga`, corresponding to what the `lspci` command shows as the vendor for the graphics chipset.

For nVidia cards, `nv` would be the free driver, but currently that one will not work in accelerated mode, and it could cause awful effects like stripes or artifacts on the screen when the *Composite* option is enabled. Sometimes these effects seem to stay forever, unless you completely disconnect the computer from the power supply and wait for a while, which seems to reset the internal memory state of nVidia cards.

If you leave out the *Driver* option, Xorg tries to autodetect the correct chipset on startup, and you get a pretty verbose protocol of Xorg initialization in `/var/log/Xorg.0.log`, which also tells you about supported compared with found chipsets.

If you have a relatively new ATI or nVidia card, the free drivers that come with Xorg might not give you the best performance or support for all resolutions. Some cards run only in generic `vesa` mode (driver `vesa` in the *Device* sec-

tion), which is very slow but still better than nothing.

If you want to run these cards with support for all their features, your only option might be to install proprietary drivers for Linux, downloaded directly from the hardware vendor's homepage. Policies and licenses for these drivers make it hard if not impossible to directly integrate them into Linux, and the work and risk of properly installing them falls on the user.

Installing nVidia Proprietary Drivers

Before starting the installer, make sure no X server is running, which you can do by switching to runlevel 2 or 3 in most distributions. Just switching to the text console with `Ctrl + Alt + F1` is not sufficient because Xorg running on a different console is still accessing the graphics hardware. First you need to stop all login managers (`xm`, `gdm`, `kdm`) that might keep Xorg running, then you can log in on the console. Installing the nVidia driver with its supplied script requires root permissions, even though not all parts of the configuration process have to run as root.

The nVidia installer is a self-unpacking and self-installing archive that you can get from the nVidia support homepage [4]. If you don't know the exact name of the card, the `lspci` command might help. If the card is very new, and it just appears as an *unknown/unlisted graphics card*, you might look at the Xorg log file in `/var/log/Xorg.log` to find out what it is.

After download, you should have a file called `NVIDIA-Linux-x86-... .run`, which you can make executable and run (as root) with:

```
chmod a+x NVIDIA*.run
./NVIDIA*.run
```

The installer is self-explanatory. nVidia's driver installation consists of a kernel module, which has to exactly match the installed kernel version and binary objects for Xorg and library functions.

The kernel module is the most breakable part of the installation. For some kernels, pre-compiled modules are present in the installer, but chances are that they won't match your installation, so the installer will try to compile a module

for your installed kernel during the installation. For this procedure to succeed, a working C compiler and development tools like the *binutils* package must be present on your system, and the kernel sources matching your installed kernel version must also be installed. If this is not the case, the installation will fail. If the necessary tools are available, a new nVidia kernel module is compiled and loaded. The installer might also add a configuration option to */etc/modules* to make sure the module is also loaded at the next startup.

A new *nvidia* Xorg driver is added to Xorg's driver module directory and configured in */etc/X11/xorg.conf*. Check for the *Driver* option in the *Device* section to see if the installer has actually activated the new driver. A backup of *xorg.conf* is created (it cannot hurt if you have made a backup on your own before running the installer). Although you cannot do much here, wait until installation is complete. nVidia provides a small utility called *nvidia-xconfig* for later reconfiguration, but the *xorg.conf* file entries for Compiz should also work with the proprietary driver (just not with the free *nv* driver).

After going back to runlevel 5 (init 5), or restarting the graphical login/session manager, the new driver module should be active. (To verify this, look for the nVidia logo shown shortly before your desktop session starts.) *glxinfo* should now show *Direct Rendering: Yes*.

Installing ATI/AMD Proprietary Drivers

Like the nVidia driver, the ATI/AMD driver requires a proprietary kernel module that matches your own kernel's version; however, the ATI installer has a graphical GUI for installation, so you won't need to shut down a running X session.

Start by downloading an installer matching your card from the AMD homepage; make it executable and run it from the command line as root:

```
chmod a+x ati-driver-2
installer-*-x86.x86_64.run
./ati-driver-installer-*-2
x86.x86_64.run
```

The graphical installer will perform all the necessary steps. Kernel source code matching your running kernel and a

working C compiler must be installed to build the kernel module.

(The ATI installer is also capable of building a Debian-based package such as those used by Ubuntu, which might be an easier way of installing the driver than the file-based method.) After installation, you can run the command *aticonfig* to reconfigure *xorg.conf*.

More Tuning

Depending on the kind of graphical applications you use, different options in */etc/X11/xorg.conf* might provide better performance, more colors, or sharper pictures. Some of the options affect all drivers, and others only work with a specific chipset driver. Options that are not known to a driver are ignored with a warning in newer Xorg versions, which is convenient because you don't have to be extra careful not to break your X server by mistyping an option.

Most options are described in detail in the online manual pages of Xorg or in Xorg's driver-specific pages. However, pay attention to the recommendations on configuring options and parameters because some graphics cards can break or overheat when "overclocked" beyond their specifications. It has been a long time since I tried a monitor with a badly chosen resolution. Most LCD displays should just cut off signals they cannot handle, but playing too much with Xorg acceleration settings might not gain enough in performance to justify the risk. Naturally, some good hints about performance tuning come from the game community, so look for useful tips and links at websites supported by Linux gamers [5]. ■

Listing 2: *xorg.conf* Options

```
01 Section "Device"
02   ...
03   Option "DamageEvents" "true"
04   Option "DisableGLXRootClipping" "true"
05   Option "XaaNoOffscreenPixmap" "true"
06   Option "AllowGLXWithComposite" "true"
07   Option "TripleBuffer" "true"
08 EndSection
09
10 Section "Extensions"
11   ?
12   # Please do NOT enable this with the "nv" driver, it may
13   # put your NVIDIA card in a confused state.
14   Option "Composite" "Enable"
15 EndSection
16
17 Section "Screen"
18   ...
19   Option "AddARGBGLXVisuals" "true"
20 EndSection
21
22 Section "ServerLayout"
23   ...
24   Option "AIGLX" "true"
25 EndSection
```

INFO

- [1] Proprietary drivers for ATI/AMD cards for GNU/Linux: <http://ati.amd.com/support/drivers/linux/linux-radeon.html>
- [2] Direct Rendering Infrastructure (DRI): <http://dri.freedesktop.org/wiki/>
- [3] Xorg homepage: <http://www.x.org/>
- [4] Proprietary drivers for nVidia cards for GNU/Linux: <http://www.nvidia.com/object/unix.html>
- [5] linux-gamers.net (with tips for Wine/Xorg configuration for various games): <http://www.linux-gamers.net/>

Protecting your system with the UFW firewall

LOCKED DOWN

Stephen Bank, 123rf.com

Careful users keep an eye on security. Ubuntu's firewall tools will help you keep intruders off your system.

BY JAMES STANGER

When it comes to security, Ubuntu offers several advantages over a system like Windows. For instance, Ubuntu's code is open source, which means hundreds and even thousands of developers review and improve the code. Also, unlike Windows, Linux systems don't have a registry, a feature that has caused Windows systems to succumb to worms, viruses, and various other attacks over the years. Linux systems are also loosely integrated, which means just because one subsystem goes bad, the rest of the system doesn't have to crash or become unstable.

Although Linux is often considered safer than other comparable desktop systems, no operating system is inher-

ently secure – not even Ubuntu. An attacker might try to scan your system to identify an open port, and the application listening on a port might have its own security issues that could lead to compromise.

But regardless of how systems or applications might compare in an abstract sense, the specific reasons really don't even matter: Configuring a personal firewall is important, and practicing safe computing helps you understand your system and brings you a bit more peace of mind.

Introducing UFW

UFW, which stands for “Uncomplicated Firewall,” is Ubuntu's official firewall configuration utility [1][2]. The com-

mand-based UFW works quite well as a front end for the powerful set of Linux firewall utilities known as iptables [3]. Although UFW appears to the user as a complete firewall system, all it really does is provide a slightly easier syntax for using the iptables.

The iptables utilities, in turn, operate the Netfilter framework found in any Linux implementation. Netfilter and iptables are complicated enough to scare off many end users, so UFW (and its GUI counterpart GUFW) were created to bring the benefits of firewalls to the everyday Ubuntu desktop user.

Getting Started with UFW

To start UFW from any standard Jaunty installation, open a terminal window

and type *ufw*. As with many of Ubuntu's administrative tools, you need to specify the root password in order to use UFW:

```
$ sudo ufw enable
Firewall is active and enabled on system startup
```

The result of the preceding command is that UFW will run each time your system starts.

After issuing this command, you can then set a default policy:

```
$ sudo ufw default deny
Default policy changed to 'deny'
(be sure to update your rules accordingly)
```

This command tells UFW to deny all ICMP packets, as well as the protocols listed in Table 1, automatically. Then you can proceed to create exceptions to this default closed policy. For example, the following commands allow connections to the SSH, web, and VNC servers running your system:

```
$ sudo ufw allow 22
Rule added
$ sudo ufw allow 80
Rule added
$ sudo ufw allow 5900
Rule added
```

You can verify UFW's status by issuing the following command:

What Is a Personal Firewall?

The term *firewall* usually refers to a device located between an internal network (such as a home or corporate network) and a larger, less controlled network such as the Internet. Although modern firewalls come with a wide range of capabilities, the most basic role of a firewall is to restrict access to the internal network by filtering incoming traffic. A personal firewall acts much like a network firewall, except that a personal firewall is typically a software-based component residing on a single system.

A personal firewall can block (or, conversely, allow) traffic based on parameters such as the source IP address, the destination address, or the port number. Services running on your Ubuntu system listen for incoming connections to a specific TCP or UDP port, so the port number provides an indication of the an outside connection is attempting to reach. See Table 1 for a summary of some common service ports.

Many personal firewall can also block or allow traffic based on the network protocol, and several firewall systems also support logging so that you can keep a record of traffic statistics and outside access attempts.

```
$ sudo ufw status
```

The output for this command is shown in Listing 1.

Allowing Specific Ports, Addresses, and Ranges

So far, the commands have only specified a port number, not an actual protocol. Filtering for the TCP or UDP protocol is also quite easy:

```
sudo ufw allow 53/udp
Rule added
```

The preceding command allows connections to your system's DNS server, which listens to UDP port 53 by default.

To allow users to connect to only TCP port 80, issue the following command:

```
sudo ufw allow 80/tcp
Rule added
```

The previous examples add rules for all IP addresses. You can get much more specific about which addresses you want to associate with your rules. For example, suppose you have a default closed policy enabled on your personal firewall and you want to allow the system at 19.82.44.45 to connect to all ports on your local system. To make this possible, you would issue the following command:

```
sudo ufw allow from 19.82.44.45
Rule added
```

Also, it is possible to specify a range of IP addresses. For example, ICANN, the body that allocates IP addresses on the

Table 1: Common Service Ports

Service Name	Description	Port Number
FTP	File Transfer Protocol server.	TCP 20 and 21 (mostly listed as 21)
POP3	The most popular email mailbox protocol. The POP3 server allows you to log in and download your messages.	TCP 110
IMAP	Internet Message Access Protocol – A newer, more sophisticated mailbox protocol that is still not as popular as POP3.	TCP 143
SMTP	Simple Mail Transfer Protocol – Email protocol used for sending and forwarding outgoing messages	TCP 25
NFS	Network File System – Used in Linux and Unix systems to share files and mount remote directories.	TCP and UDP port 2049; also relies on the Portmapper service, which uses UDP port 111
SSH	A relatively secure way to access and control a system, usually via the command line. Uses encryption and public-key-based authentication. You can also use SSH to tunnel unencrypted protocols.	TCP 22
Telnet	An older protocol (mostly replaced by SSH) used to control remote systems.	Server listens on TCP 23; you can use a Telnet client to connect to any port
VNC	Virtual Network Computing – Protocol that allows you to log in to a remote computer and view the remote X Window session as if you were sitting in front of it.	TCP 5900, by default
IPP	Internet Printing Protocol – Used to connect printers across networks.	UDP and TCP 631.

Default Policy

You can configure a personal firewall to take one of the following default configurations:

- Default open
- Default closed

UFW allows you to choose the default stance. I follow conventional wisdom and choose the default closed configuration, because it's generally better to disallow everything and then permit only specific traffic. The advantage to this policy is that you have granular control over the types of connections your system allows.

The disadvantage is that you might inadvertently block connections to your system and will have to spend time creating exceptions to this default policy. But as you will learn in this article, creating these exceptions is relatively easy.

Listing 1: ufw Status Output

```
01 Status: active
02
03
04
05 To Action From
06
07 -- -----
08
09 22 ALLOW Anywhere
10
11 80 ALLOW Anywhere
12
13 5900 ALLOW Anywhere
```

Internet, created some ranges of private IP addresses that can't be routed on the Internet. These addresses are often assigned by DHCP servers operating behind a firewall.

To allow all of the RFC 1918 "private IP" addresses to access your system, you would issue the following commands:

```
$ sudo ufw allow from 10.0.0.0/8
rule added
$ sudo ufw allow from 172.16.0.0/12
rule added
$ sudo ufw allow from 192.168.0.0/16
rule added
```

Of course, you can specify any IP address range you like.

Limiting Connections

One of the more powerful features of UFW is the ability to limit connections.

Won't It Block All Access?

Many users won't enable a personal firewall because they're afraid they then won't be able to visit their favorite video site or download programs. That's really not going to happen. The developers who create personal firewall software are very smart; they know that you want to get out to all the good stuff on the Internet. They also know there are plenty of bad guys who are trying to get into your computer.

The personal firewall software generally blocks incoming access to your system; it won't hamper your ability to go out and find what you want. Unless, of course, you specifically tell it to.

For instance, you can specify that your UFW personal firewall system will allow no more than six connection attempts over a period of 30 seconds.

Limiting the number of connection attempts can help you thwart the following types of attacks:

- Scanning: If connections are limited, scans are less accurate.
- Denial of service: A denial of service attack involves sending floods of packets or malformed packets to a host in the attempt to either crash the system or overwhelm it. Limiting connections can help the system ignore flooding attempts.
- Brute force: A brute force attack is where an attacker repeatedly tries to guess a user name or password combination. Limiting connections causes a reset to occur, causing the attacker to take far longer to find a successful username/password combination.

To enable connection limiting, use the following syntax:

```
ufw limit service_name/ protocol
```

For example, to limit connections for your web server, you could issue the following command:

```
$ sudo ufw limit www/tcp
rule added
```

Removing a Rule

If you want to remove a rule, specify the port. For exam-

ple, to remove the rule allowing SSH, issue the following command:

```
$sudo ufw delete allow 22
Rule deleted
```

If you want to change the default stance of your firewall, simply issue the following command:

```
$ sudo ufw default allow
```

However, understand that all of the existing rules you have created will still exist, so it is best if you stick with a default deny policy, then create rules to allow specific traffic.

GUFW

Operating your firewall from the command line gives you the most control over your personal firewall, but you don't have to use the command line if

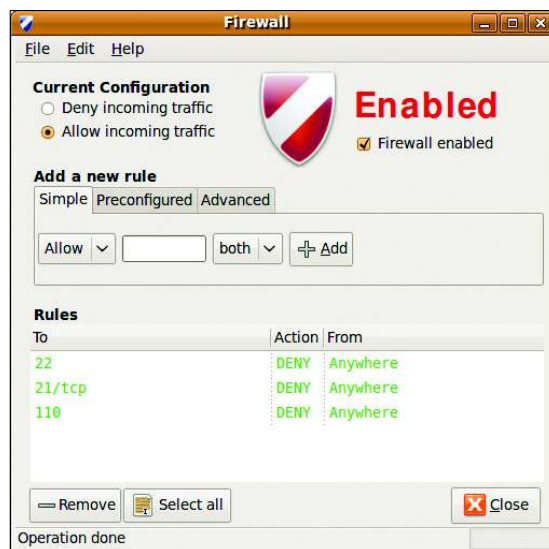


Figure 1: UFW, showing a simple firewall configuration.

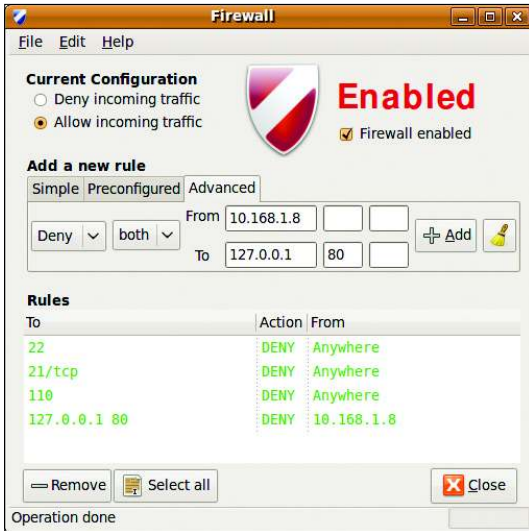


Figure 2: Specifying advanced port features in GUPFW.

you don't want to. Jaunty provides a fairly robust GUI application known as GUPFW (Figure 1). To access GUPFW, go to *System | Administration* in the Ubuntu main window, then select *Firewall configuration*. Just as with UFW, you will need to provide your root password to launch GUPFW.

Getting Started with GUPFW

Once you've launched GUPFW, you will see that three tabs are available:

- Simple: Allows you to add a particular port number or service quickly.
- Preconfigured: Instead of requiring you to specify a service, GUPFW provides a ready-made list for you.
- Advanced: Allows the creation of more sophisticated rules, including source and destination ports and IP addresses.

Figure 2 shows how you can add an advanced rule that denies both TCP and UDP traffic from the system at IP address 10.168.1.8 to the local system.

Logging

Logging is a vital function of any personal firewall. To enable logging in

Sidebar: Installing UFW and GUPFW

If, for some reason, UFW and GUPFW aren't installed on your Ubuntu system by default, make sure that the Universe repository is enabled, and then issue the following commands:

```
sudo apt-get install ufw
sudo apt-get install gufw
```

GUPFW, go to *Edit | Preferences*, and then select the *Enable ufw logging* checkbox.

To enable logging from the command line, enter:

```
ufw logging level
```

If you use the GUPFW or the *ufw* command to turn logging on without any arguments, the logging level will automatically be set to low. To set the level to full (the highest level of logging), issue the following command:

```
$ sudo ufw logging full
logging enabled
```

Specifying the full or even the *high* setting will cause your system to log an enormous amount of connections, which might fill up your hard drive. I recommend specifying low logging, which is the default.

Logging isn't just an all or nothing proposition. You might want to log only specific connections. The following command tells *ufw* to log connections to your VNC server (port 5900):

```
ufw allow log 5900/tcp
```

The following command disables all logging:

```
$ sudo ufw logging off
logging disabled
```

Some Final Tips

Most users configure GUPFW to start when they log on to the GUI desktop system. To ensure that this happens,

Other Firewalls

If you don't want to use UFW, you don't have to. Other firewall creation utilities have existed for years, including IPCop [4], SmoothWall [5], Endian [6], or Arno-iptables [7]. Some of these products are pretty robust and provide more features than a typical end user really needs. Still, it's good to know about these tools in case you ever need to use them in the future.

start GUPFW, then go to *Edit | Preferences* to open the Preferences dialog box, shown in Figure 3.

Under the *System Settings* section, select all three checkboxes, as shown. Next time, when you boot your system and log in, GUPFW will be running automatically, nicely tucked away in your status bar.

Once you have created a rule set that you like, you might want to save a backup. To do so, go to *File | Export rules*, and then save the configuration as a simple text file. Then you can load this configuration back into your system – or onto another system – by going to *File | Import rules*.

Conclusion

My advice to you is to play around with this firewall a bit. One person's perfect

firewall is another's perfect nightmare. I guess that's why they call it a "personal firewall."

If I could propose a one-size-fits-all solution, I'd offer up a text file you could just import into the firewall, and you'd be finished.

But it really doesn't work that way for a couple of reasons. First of all,

your needs are different than mine. Second, it's important for you to take responsibility over your computer and customize it for your needs. If you work on learning these settings, we'll all be one step closer to a safer Internet. ■




Figure 3: Setting up GUPFW to run automatically in Jaunty.

INFO

- [1] The Ubuntu UFW Wiki: <https://wiki.ubuntu.com/UbuntuFirewall>
- [2] The Ubuntu UFW Forum: <http://ubuntuforums.org/showthread.php?t=823741>
- [3] Netfilter/iptables: <http://www.netfilter.org/>
- [4] The GUPFW home page: <http://gufw.tuxfamily.org/index.html>
- [5] IPCop: <http://www.ipcop.org>
- [6] SmoothWall: <http://www.smoothwall.org>
- [7] Endian: <http://www.endian.com>
- [8] Arno-iptables: http://freshmeat.net/projects/iptables-firewall/?topic_id=151

Running Windows programs on Ubuntu with Wine

WINE IMPORTER



Shanna Korby, Fotolia

Users who move from Windows to Ubuntu often miss some of their favorite programs and games. Wouldn't it be practical to run Windows applications on the free Ubuntu operating system? Time for a little taste of Wine. **BY TIM SCHÜRMANN**

Many Ubuntu migrants miss games and graphics programs such as CorelDRAW or products such as Adobe Photoshop. The only solution is to install Windows parallel to Ubuntu – or try Wine, which tricks applications into believing they are running on a Windows system.

The history of Wine goes back to the year 1993. At the time, Sun developed a small tool to run Windows applications on its own Solaris operating system, without the need for a Windows license or other interaction. This software inspired Bob Amstadt and Eric Youngdale

to develop something similar for Linux. A short while later, the first version of Wine was released. Today, more than 300 volunteer programmers from all over the world continue to contribute to the Wine project.

What's in a Name?

To run Windows programs on Ubuntu, Wine uses a fairly complex trick: It sits between the Windows application and Ubuntu like a simultaneous interpreter. If the Windows application says: "Draw a button," Wine passes this command directly to Ubuntu. In contrast to Virtual-

Box or VMware, Wine does not emulate a whole PC and thus cannot be considered a real emulator. This also explains the name Wine, which means Wine Is Not an Emulator.

Because of the way Wine works, it offers a number of advantages. Chiefly, you do not need an expensive Windows license. Programs will run almost as fast as on the Redmond operating system, and windows behave as if they belong to a native Ubuntu program.

An interpreter can only be as good as its knowledge of the two languages, and therein lies the problem: Because Micro-

Version	Description	Rating	Wine version	Test results	Comments
10	Excellent vector graphic editor, best on market.			0	20
11	CorelDRAW Graphics Suite composed by: CorelDRAW, PHOTO-PAINT and CAPTURE. More than 10,000 individually selected clipart and digital images, and 1,000 OpenType fonts.	Silver	0.9.51.	16	35
12	Vector programa	Garbage	1.1.5	15	9
7.0 - Select Edition	This version is a suite to small offices.	Gold	1.0-rc4	5	0
9	CorelDRAW is a popular vector	Gold	0.9.15.	1	28
X3 (13)	CorelDRAW X3 Graphics suite (version 13). Vector and bitmap drawing suite.	Silver	1.1.14	12	6
X4	CorelDraw X4 Graphics Suit includes:	Garbage	1.1.5	5	0

Figure 1: The Wine database lists numerous applications that run on Wine - for example, the CorelDRAW graphics package.

softs refused to release details of its operating system until just recently, Wine developers were forced to adopt a trial and error approach. Such a method involves time-consuming, painstaking work, and considering the multiple-gigabyte operating system monster they were facing, it was very much a labor of Sisyphus. On top of this, some software manufacturers use some pretty nasty programming tricks, and if you want to talk these applications into running on Ubuntu, Wine needs to know these tricks. Unfortunately, most software vendors keep this kind of information safely under lock and key.

Despite years of work, then, Wine is still officially under development, and most of the work taking place now is going into the fields of multimedia, video, and 3D graphics. This said, Wine does support DirectX 8 and 9 quite well by now. The developers started work on the tenth incarnation of this critical graphics interface for games last year. An additional limitation is that Wine can only handle very simple 32-bit Windows applications, although improved support for 64-bit software is being worked on even now.

Wine has another disadvantage: It can only handle Windows applications; it

cannot revitalize hardware with Windows-only driver support.

Preparation

Before you fire up your package manager and launch into the Wine install, you might want to take a look at the application database on the Wine homepage [1]. This database summarizes the level of support for various Windows programs: whether or not they run on Wine and, if so, how well behaved they are. To access the database, go to the Wine homepage and, from among the tabs at the top, click *AppDB*. In the *AppDB* area on the left, select *Browse Apps*, go through the drop-down list labeled *Category* until you find the category and application you are interested in, and click on *Update filter*. CorelDRAW and Photoshop are hidden below *Multimedia | Graphics | Graphics Editing*, for example. Alternatively, use the search function to find your application in the database.

If you decide to investigate CorelDRAW, you are taken to the page shown in Figure 1, where you will find ratings for different versions that tell you how well the

program fared in tests performed by Wine users. The *Platinum* rating is only awarded if a program runs without errors and without any impairment to functionality on Wine. A couple of older CorelDRAW versions have made *Gold* status; the current X4 version refuses to cooperate entirely (*Garbage*). Adobe Photoshop CS2 will run more or less as the makers intended, whereas the current CS3 version has reached only *Bronze* status. Things are even worse in the case of Adobe Illustrator. As a rule of thumb, the older a program is, the better your chances are of getting it to work.

A Wrench in the Works

Clicking on a program version tells you exactly what the trouble is. A new window appears with details of what does and doesn't perform as expected. In addition, user reports at the bottom and tips for talking the program into cooperating can be very useful. The steps this requires, some of which are very complex, are something you should not attempt unless you are an experi-

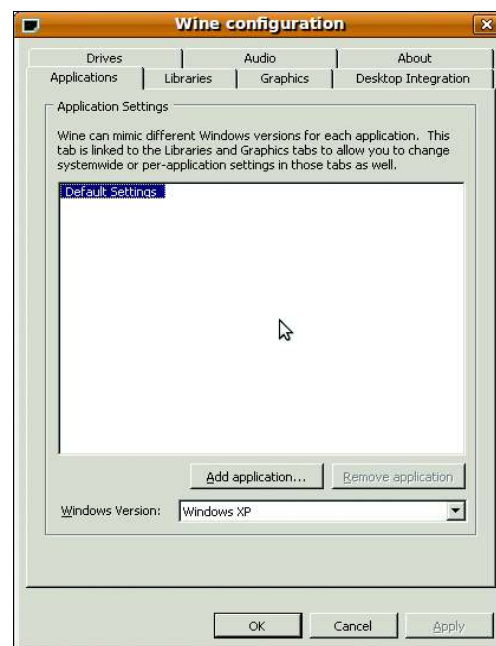


Figure 2: This window shows you the current Wine settings.




Figure 3: Start the Windows application installation file using Wine.

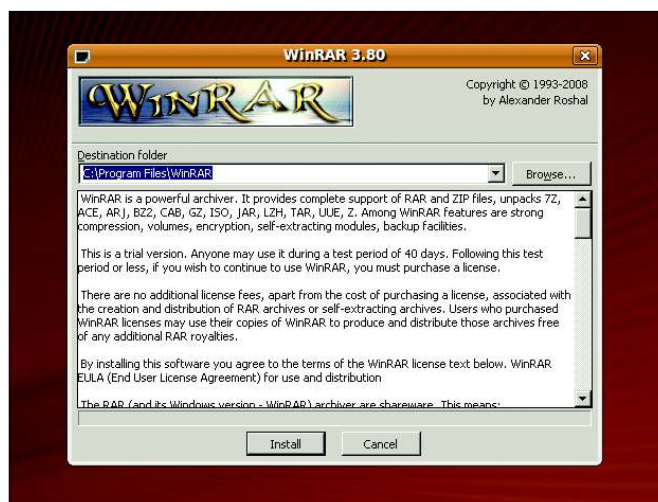


Figure 5: Installing WinRAR with Wine.

enced Ubuntu user. In some cases, programs refuse to work because Internet Explorer is missing; for example, Lexware accounting software relies on Internet Explorer for some operations.

If your application is still listed as broken in the database, you should check the test date (in *Test Results*). If the test was a while back, you might want to experiment yourself. Wine continues to mature, and it is conceivable that your application has been taken care of.

Once you have established that your Windows application is likely to run on Wine, you need an installation disc for the Windows application. Although Wine is free of charge, you need a valid license for the Windows application.

Installation

Wine is included with Ubuntu and you can retrieve it from a repository. First,

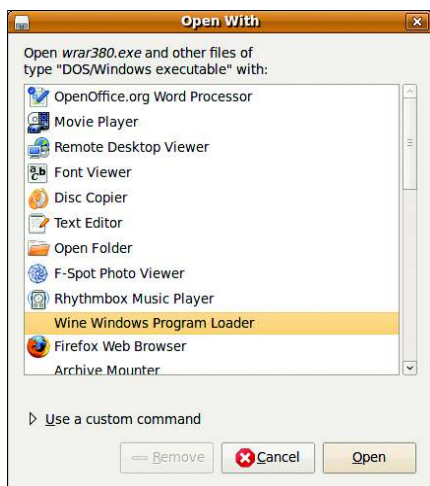


Figure 4: Choose "Open With" and select "Wine Windows Program Loader" for Windows applications.

Launch *System Administration Synaptic Package Manager*, enter your password in the resulting dialog, then type *wine* in the search box and click *Search*. In the list of results, mark the *wine* entry for installation. Synaptic will present a list of additionally required changes. Now confirm these with *Mark* and start the installation by clicking *Apply*.

Preparations

Wine integrates with the desktop start menu. To launch Wine, click on *Applications Wine Configure Wine*.

The Wine Configuration tool will then go on to create a hidden *.wine* directory below your home directory. The software uses this directory to store the basic configuration, which you can see in Figure 2. Right now, do not modify anything here, but click *Cancel* to close the tool.

Installing a Windows Program

To take Wine for a test run, download the WinRAR packer [2], then open the file manager. Next, locate the file and

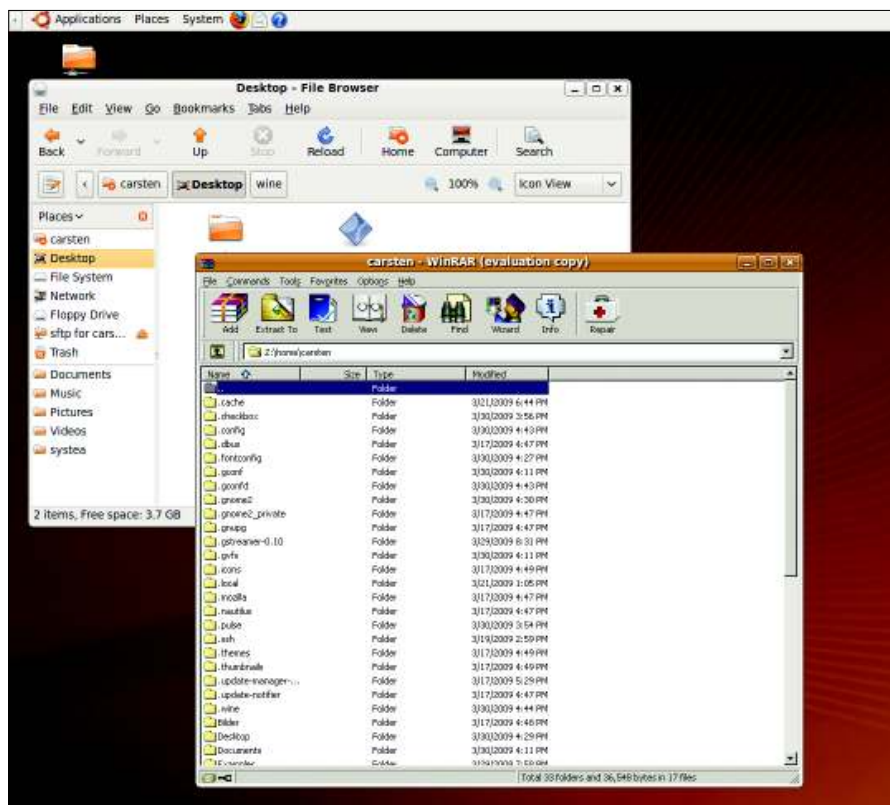


Figure 6: The Windows application WinRAR running on the Ubuntu desktop.

start the installer by double-clicking on it. Alternatively, choose *Open with "Wine Windows Program Loader"* from the context menu, which opens after a right-click on any file (Figure 3).

In WinRAR's case, the self-extracting file *wrar380.exe* automatically executes the installer. In other cases, you might need to look for the installer for the Windows application in an archive file: It will typically be *setup.exe*, or *autorun.exe*.

If Wine does not launch automatically after double-clicking on the installer, you should choose *Open with Other Application...* (Figure 3), then choose *Wine Windows Program Loader* in the window shown in Figure 4. Now you just need to click *Open*.

Wine takes over in the background and launches the Windows program. Just follow the normal steps to install the program. Although I talk about Windows drive letters in the next section, for now, confirm the default installation directory suggested by the application (Figure 5). If any installer asks you to reboot Windows, press Alt+F2, type *wineboot*, and click *Run*. This tells Wine to simulate a Windows reboot. Figure 6 shows how a Windows program looks running in Wine. If Wine does not launch the Windows program, check out the "Dead as a Dodo" box.

If the installation program creates one or multiple entries in the start menu, Ubuntu places them below *Applications | Wine | Programs*. If so, you can use the

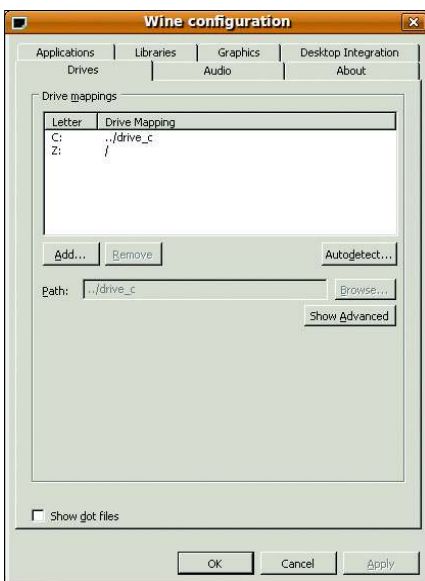


Figure 7: The Drives tab lets you map Linux directories to drive letters.



Figure 8: Wine has a graceful approach to uninstalling Windows programs.

entries as a convenient method for launching the Windows program.

Tell Me Your Name

By default, Wine stores Windows programs that think they are installing on drive C: in a hidden subdirectory below your home directory. To display the hidden files, select *View | Show Hidden Files* in the Nautilus file browser, then change directory to *.wine* in your home directory. One level below is a *drive_c* subdirectory and below this, *Program Files* and *Windows*. The former contains the application you just installed. After the WinRAR installation, it will contain a folder with this name with all the application files. To start WinRAR, double-click the file *WinRAR.exe* in its folder. Also, you can run Windows applications from a terminal window by entering *wine* and the executable's file name.

Directory Jungle

Unlike Windows, Ubuntu does not use drive letters, so Wine must use a work-

around: Each drive letter maps to an Ubuntu directory. Although a Windows program running on Wine thinks it is storing data on drive C:, it is writing to a subdirectory called *.wine/drive_c* below the home directory. The Wine control center shows which letter maps to which directory. First choose *Applications | Wine | Configure Wine* or type *winecfg* in a quick launcher (Alt+F2). Then visit the *Drives* tab (Figure 7). The list at the top shows you which drive letter maps to which Ubuntu directory. Normally, C: maps to *.wine/drive_c*, and Z: lets the Windows application access the whole directory tree because it maps to the root directory. Here, you can *Add* mappings by defining new virtual drive letters and typing the *Path* to which the letter should point.

VistaXP 98

Some Windows programs require the characteristics of a specific version of Windows. The *Applications* tab can help you cater to this.

Dead as a Dodo

If a Windows program fails to launch when you click it, some manual attention might make it see reason. To start, open a terminal window. Then, select *Application | Accessories | Terminal* in the start menu. Now type

```
cd directory
```

where *directory* represents the full path to your installation program. If the *setup.exe* file is on a DVD, you can access the DVD via the */media/path*. The command line looks like this:

```
cd /media/DVD_name
```

In this command, replace *DVD_name* with the entry in your *media* directory, then climb down the directory tree until you find the directory with the *.exe* file you

need (in WinRAR's case, this is */programs/other/winrar*). When you get there, type *wine file*

replacing *file* with the name of the application you want to run, say, *setup.exe*. Make sure you stick to the exact spelling, remembering that upper- and lowercase letters are vital. To tell Wine to run the Windows program, press Enter again.

If you experience difficulty launching the program, check the terminal window for error messages – you should not close the terminal window or you will terminate the Windows program running in it. If the Windows program crashes, you can try relaunching it.

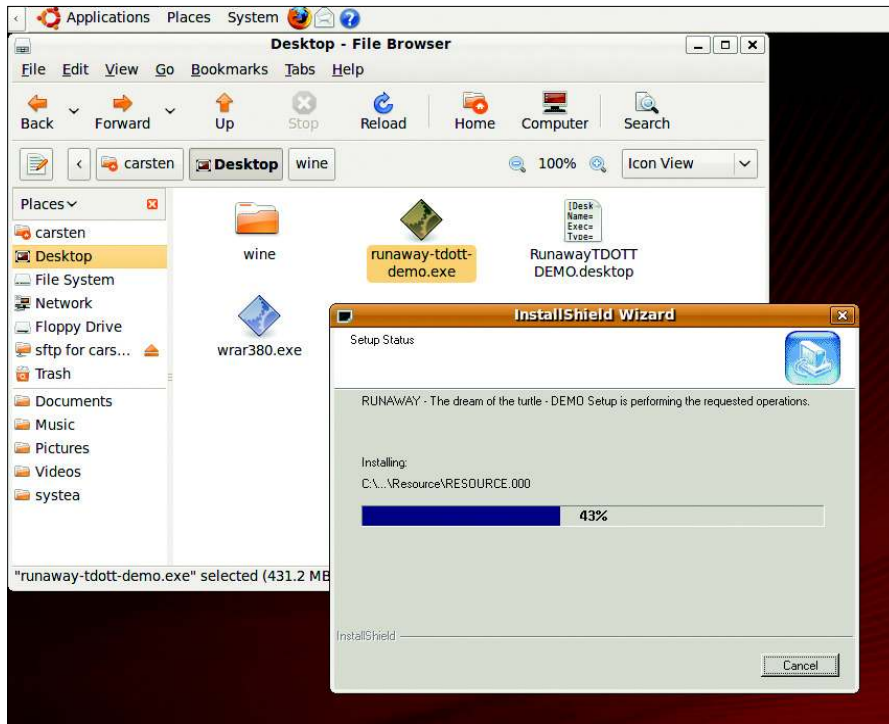


Figure 9: Some Windows games, such as Adventure Runaway, ...

For each application, you can individually define the Windows version that Wine should emulate. To do so, click *Add application* and look for the executable for your application (in `.wine/drive_c/`).

Now click *Open*, select the application in the list, and then select the required version of Windows in the *Windows Version* drop-down list. The configuration in *Default Settings* is valid for all Windows

applications, unless you define an exception.

Garbage Collection

To get rid of a Windows application you have installed on Ubuntu, you might think it would suffice to delete the corresponding directory below `.wine/drive_c/Program Files`.

However, Wine will still list the application as installed. Instead, run *Applica-*

tions | Wine | *Uninstall Wine Software* or press `Alt+F2` and type `uninstaller`. Now you should see the window in Figure 8. When you get there, select the software you want to delete and click *Uninstall*.

Conclusions

Wine is not a panacea, and many major programs still refuse to cooperate. Your best chances are with games (see Figures 9 and 10), but you might need to do without popular accounting packages, CorelDRAW, the latest version of Photoshop, or iTunes, for example.

A better alternative might be to look around for an Ubuntu program that provides similar functionality. In the case of Photoshop, why not try GIMP or Krita? And OpenOffice has long been a more than adequate replacement for Microsoft's Office suite. If you check out the list of packages in your distribution's package manager, you might be surprised at the number of comparable programs, or even programs with superior features, that Ubuntu offers. ■

Apps Gone Bad

The Wine MIT license gives users many privileges, a fact that has helped TransGaming generate huge profits. TransGaming just grabbed the package, added DirectX capabilities, which are required for gaming, and started to sell the results as Cedega [3] (formerly WineX). The Wine developers were unhappy about this, and thus changed the licensing terms to the LGPL in March 2002. The LGPL ensures that you can continue to use Wine for commercial purposes, but that changes and additions must be given back to the project. CodeWeavers is exemplary in this respect; the company bundled Wine with its own commercial components to create the CrossOver [4] package, while actively promoting the ongoing development of Wine.

INFO

- [1] Wine Project homepage: <http://www.winehq.org>
- [2] WinRAR archiver homepage: <http://www.rarsoft.com>
- [3] Cedega website: <http://www.cedega.com>
- [4] CrossOver Linux: <http://www.codeweavers.com/products/cxlinux/>

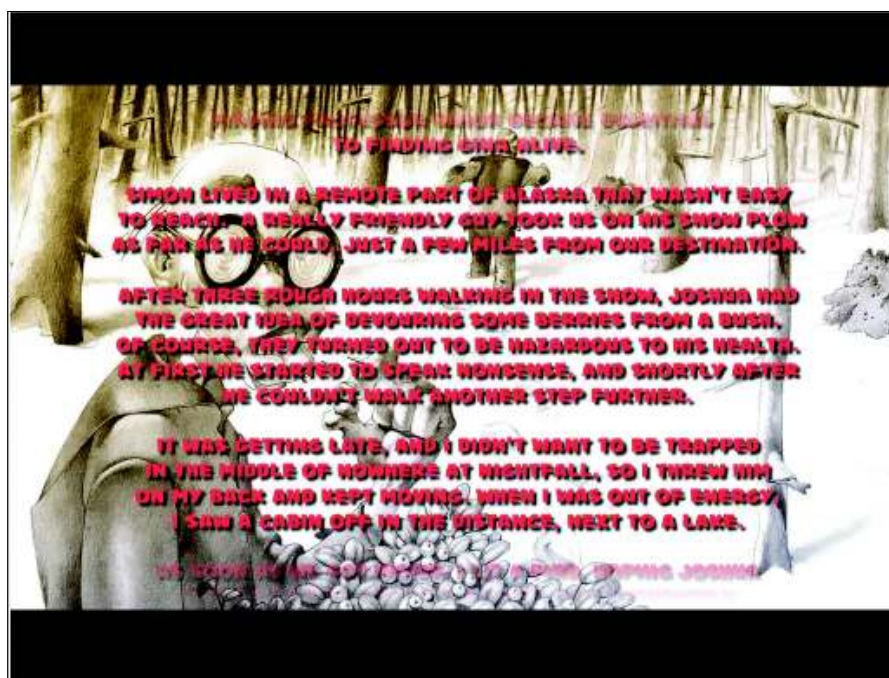


Figure 10: ... run on Linux thanks to Wine.


Expert Touch

Linux professionals stay productive at the Bash command line – and you can too. The Linux Shell special edition provides hands-on, how-to discussions of more than 300 command-line utilities for networking, troubleshooting, configuring, and managing Linux systems. Let this comprehensive reference be your guide for building a deeper understanding of the Linux shell environment.

You'll learn how to:

- Filter and isolate text
- Install software from the command line
- Monitor and manage processes
- Configure devices, disks, filesystems, and user accounts
- Troubleshoot network connections
- Schedule recurring tasks
- Create simple Bash scripts to save time and extend your environment

The best way to stay in touch with your system is through the fast, versatile, and powerful Bash shell. Keep this handy command reference close to your desk, and learn to work like the experts.



AVAILABLE AT YOUR NEWSSTAND!
OR ORDER ONLINE AT WWW.LINUXPROMAGAZINE.COM/SPECIAL



Tools and techniques for building a website

NEW PAGES

With readily available Ubuntu tools, you can create and post your personal websites to a server. We'll show you how to get started. **BY JAMES STANGER**

Sophisticated corporate sites include content management tools, server-side scripting, and back-end database systems, but if you just

want to start building your own website, the free tools available through the Ubuntu repositories provide all you need to create web pages, podcasts, graphics,

and all of the other elements of a solid website.

Even if you're a beginner, Ubuntu's Jaunty Jackalope supports the applica-

HTML Editors

Cream

Cream is a terrific little text editor. Based on vim, Cream is an X Window application that makes it easier for people who want to get right into the code. If you can use any standard Windows application, you can use Cream.

To install it, make sure you have the standard Jaunty Universe repositories enabled and type the following:

```
sudo apt-get install cream
```

Cream only takes up about 2.5MB of disk space on my standard Jaunty system.

Firefox Add-ons

The Firefox web browser can serve as a sophisticated development environment, as long as you install the correct add-ons. Some of my favorites are:

- Firebug [1]: Turns Firefox into an HTML editor.
- Web Developer [2]: Inspects all web page

elements, conducts validation, and performs many other functions.

- YSlow [3]: Shows you exactly how long your page will take to download.
- SenSEO [4]: Gives you hints and tips for making sure the pages you create are optimized for the search engines.

If you use Firebug, you get the best of both worlds: a strong text editor and a quasi-WYSIWYG editor. Also, you really don't need to install any other applications.

Bluefish

If you find that using a straight text editor such as Cream is too technical, Bluefish [5] is a nice alternative. Bluefish (Figure 2) is a full-featured HTML editor that provides features such as spell checking, linking to style sheets, and automatically launching various browser types for validation.

Quanta Plus

Quanta Plus (Figure 3) [6] is a KDE-based application that offers image maps, CVS

uploads, link validation, and other sophisticated features in addition to basic HTML editing. The user-friendly Quanta Plus organizes HTML code in an intuitive way. For instance, you can select tags from a convenient drop-down menu.

Kompozer

If you just want a WYSIWYG editor that operates more or less like a word processor, the best option might be Kompozer [7]. Formerly known as NVU, Kompozer is easy to use, and it sports some relatively sophisticated features, such as a CSS editor and JavaScript console.

Also, you can put Kompozer into "code mode," wherein the application appears much like a text editor. An intermediate "tag mode" is also available. Tag mode is helpful because it shows you where each of the HTML tags occur on a page. Double-clicking on a tag allows you to edit the code.




Figure 1: The directory structure of a typical web development environment.

tions you need to be creative and build a personal web presence.

The Big Picture

The task of developing a website generally includes the following steps:

1. Create the HTML from a template.
2. Add images.
3. Test, review, and validate your work.
4. Post your work on a staging server, and then on the production server.

The web page that appears in your browser as a single, unified image might actually consist of several files, including digital images, scripts, and Hypertext Markup Language (HTML) files.

As you probably know, HTML is the markup language at the center of the World Wide Web. An HTML file is a text file with a set of additional codes called *tags* that define how the file will appear in a browser window. An HTML tag can

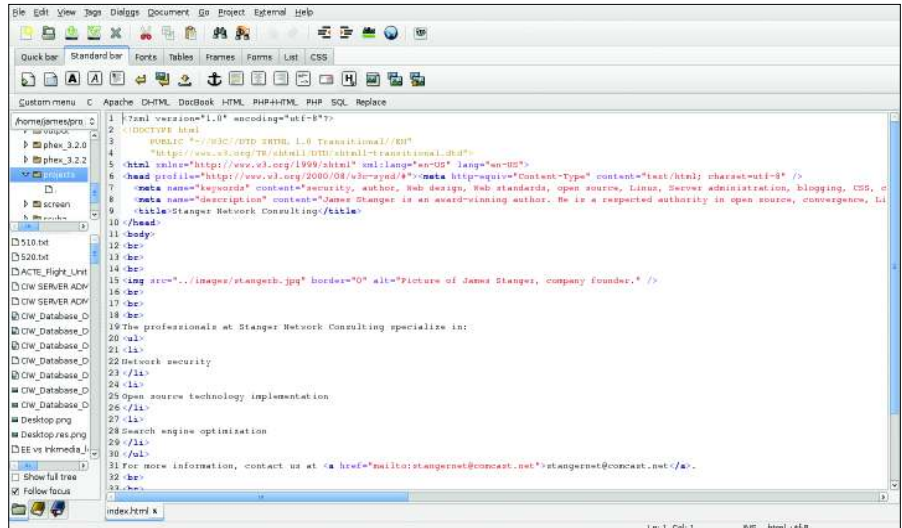


Figure 2: Editing an HTML file in Bluefish.

specify a font, a link, an image, or an instruction related to the page layout.

A full tutorial just on HTML would require a whole book. Several excellent HTML tutorials can be found online. For the purposes of this article, though, keep in mind that you no longer have to hand-code HTML tags into text files. Ubuntu supports several web editing tools that let you work with the page as it will appear to the viewer – without getting hung up in the details of HTML.

The box titled “HTML Editors” describes some editing tools available for the Ubuntu environment. These tools serve the role commonly associated with tools such as Adobe DreamWeaver in Windows and Macintosh environments, except most of Ubuntu’s tools are available for free through the Ubuntu repositories. For more on finding and installing software in Ubuntu,

see the article on package management in the Discovery guide.

Establishing a Development Environment

As you begin building your website, it is a good idea to create separate directories for your HTML pages, images, and scripts (Figure 1). Organizing your files into a discrete, well-defined directory structure helps ensure that your pages will render properly in a web browser after they’ve been transferred to a web server system.

Rather than working on a live Internet site, most developers do their web development and testing on a staging server that is as close as possible to the final web server environment. Working on the staging server lets you review and debug code in a nearly real environment, but if you make a mistake, the general public

Listing 1: HTML Template

```

01 <!DOCTYPE html
02     PUBLIC "-//W3C//DTD XHTML
03     1.0 Transitional//EN"
04     "http://www.w3.org/TR/
05     xhtml1/DTD/xhtml1-transitional.
06     dtd">
07 <html xmlns="http://www.
08     w3.org/1999/xhtml" xml:lang="en">
09 <head>
10 <meta name="Keywords" content="CIW,
11     Foundations, Example"/>
12 <meta name="Description"
13     content="For the CIW Foundations
14     Course"/>
15 <meta http-equiv="Content-Type"
16     content="text/html;
17     charset=utf-8"/>
18 <title>XHTML Template</title>
19 </head>
20 <body>
21 <p>This is an XHTML file. It validates
22     to the XHTML 1.0 Transitional
23     standard.
24 </body>
25 </html>
  
```

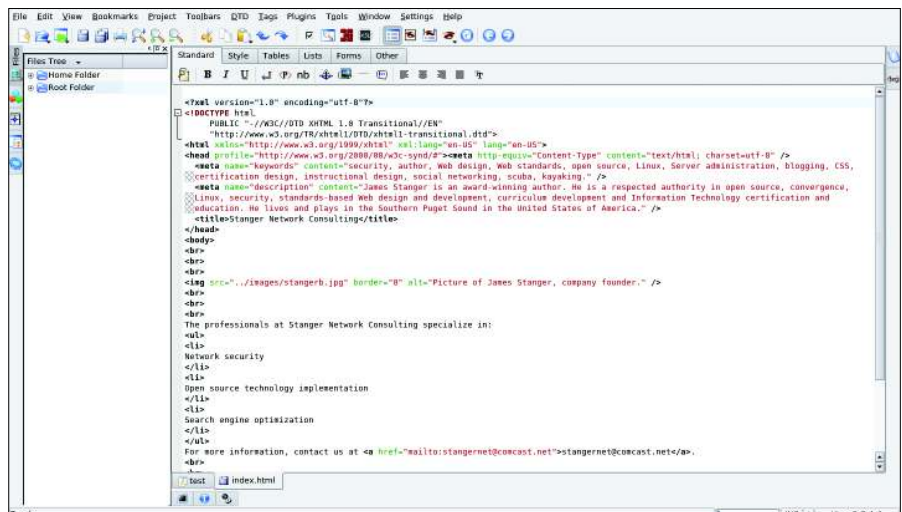


Figure 3: Editing an XHTML page in Quanta Plus.

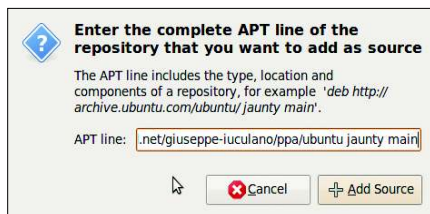


Figure 4: Entering a repository in Synaptic. If the APT line is too long for the box, just keep typing.

won't see it. Once you are satisfied with the web pages, you can upload the directories to a production server.

Templates and HTML

Whenever I create a web page, I always start with a standard template file. The template file contains standard information that allows me to create a consistent set of pages. Listing 1 shows a bare-bones HTML template. As you can see, the template provides a standard form for calling out parameters such as the HTML version, the language, and the character set. Other statements let you designate keywords and descriptive text that will be associated with the site.

The very simple template shown in Listing 1 provides only a fraction of the possible options, but it is a good beginning. If you are looking for a more sophisticated alternative, try the Open Source Web Design [8] or Open Source Templates [9] websites.

The `!DOCTYPE` statement, which specifies the HTML standard, is an important element of the template. Often I'm asked which version of HTML or XHTML people should use. Although I prefer XHTML 1.0 Transitional – it is supported by the majority of web browsers, is perfect for data-oriented social networking applications, and is relatively current – to be honest, you can choose any version of HTML you want,




Figure 5: Kompozer opens in WYSIWYG mode.

including good old HTML 4.01. Just pick a standard and use it consistently.

Because an HTML file is essentially a text file, you can create and edit HTML files with an ordinary text editor, such as vim, emacs, or the gedit editor available through Ubuntu's user interface (*Accessories | Text Editor*). A pure text editor, however, doesn't check your code or provide any particular coding advice.

An HTML editor is specifically designed for the task of creating web pages. HTML editors inform you about the quality of the code you're entering. An HTML editor will also make it easy for you to insert standard HTML tags.

Several HTML editors are available for Ubuntu, and the ideal choice is often a matter of personal taste (refer to the box labeled "Choosing an Editor").

Kompozer

For this article, I'll focus on setting up a page using Kompozer. The other editors are similar – see the documentation for your own editing tool. Kompozer doesn't appear in the default Ubuntu user inter-

face, so you'll have to add it. If you're comfortable working from the command line, you can use the apt-get utility. If not, you can add Kompozer through the Synaptic package manager.

To open Synaptic, choose *System | Administration | Synaptic Package Manager* in the Ubuntu main window. To access Synaptic, you'll need to enter the administrative password (the password you defined when you installed the system).

In the Synaptic main window, choose *Settings | Repositories* and select the tab labeled *Third-Party Software*. In the *Software Sources* dialog box, click the *Add* button. Synaptic will prompt you to enter the APT line. Enter the following (as shown in Figure 4):

```
deb http://ppa.launchpad.net/ &
giuseppe-iuculano/ppa/ &
ubuntu jaunty main
```

Next, click the *Add Source* button to add the repository, then close the *Software Sources* dialog box and click *Reload* in the Synaptic main window. Synaptic will

Choosing an Editor

A few important qualities to look for in an editing tool include:

- **Standards compliance:** An HTML editor should support open standards and make it easy to choose a standard, from HTML 4.01 to XHTML 1.0 Strict and beyond.
- **Cascading Style Sheets (CSS) support:** CSS is a method for adding style elements to web pages. As you advance to more sophisticated web development assignments, your job will be much easier if you're using an editor that supports CSS.
- **File publishing:** A good HTML editor should provide a means for uploading web pages to the server with a minimum of fuss.
- **Ease of use:** An application that gets you up and running quickly might be worth the sacrifice of multiple feature sets.

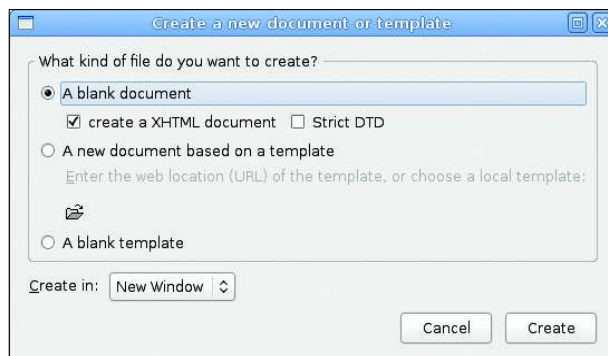


Figure 6: Creating a new document in Kompozer.




Figure 7: Creating a bulleted list in Kompozer.

add the contents of the repository to the package list.

If the reload is successful, enter *kompozer* in the *Quick Search* box. When *kompozer* appears in the package list, select the checkbox to add the package and click the *Apply* button. Synaptic will add Kompozer to your Ubuntu system. (For more information on adding new applications to your Ubuntu system, see the article on adding packages in the Discovery guide.)

Lauching

On standard Ubuntu systems, Kompozer will not appear automatically in the menu once you add it. To start Kompozer, launch a terminal window by selecting *Applications | Accessories | Terminal* in the Ubuntu main window and type *kompozer* (all lowercase – remember that Ubuntu is case sensitive) at the command prompt. Kompozer starts with the user-friendly interface shown in Figure 5. *Note*: Some early versions of Jaunty have experienced issues with Kompozer. These issues will likely be

Spell Checking

Right now, Kompozer doesn't have a good spell checking feature. If you're stumped for decent spell checking in Kompozer, consider copying and pasting text from Kompozer into OpenOffice.

cleared up by the time you read this, but the Ubuntu forum has posted some information for early adopters [10].

Creating a Document and Formatting Text

To begin a new document in Kompozer, choose *File | New*. Now you'll be presented with the *Create a new document or template* box. To create an XHTML 1.0 Transitional document, edit the dialog box so that it looks similar to the box shown in Figure 6.

Once you click *Create*, you'll then be presented with a blank document. Just start typing inside the document as you would any word processor. As you type

in some text, you can also apply formatting, including bullet points, headers, and so forth. To add bullets, create three new lines of text, then highlight them and go to *Format | List | Bulleted*.

The list you've just created is now bulleted, as shown in Figure 7.

As you explore the Kompozer menus, you will find other options for aligning, resizing, and italicizing text.

Adding an Image

Most web pages include graphic images as well as plain text. With the use of an image processing tool, you can add a photo image or create a graphic element for the site. GIMP is a popular graphics tool that is accessible through the Ubuntu Applications menu (*Applications | Graphics | GIMP Image Editor*). For a quick look at how to build a simple graphic header for your website, see the sidebar titled "Creating a Simple Header Bar."

To add an image in Kompozer, place your cursor in the document where you want the image to reside. For example, place the cursor near the top of the page. Then, choose *Insert | Image*, and browse to an image you have previously created and saved.

Now enter the location of the image and some alternative text that describes it. The alternative text helps search engine bots organize the information on the site, which could increase the search rank for the page. Also, screen reading tools use the alternative text to describe the contents of the image to visually impaired users. Once you have inserted the image, it will appear in the Kompozer

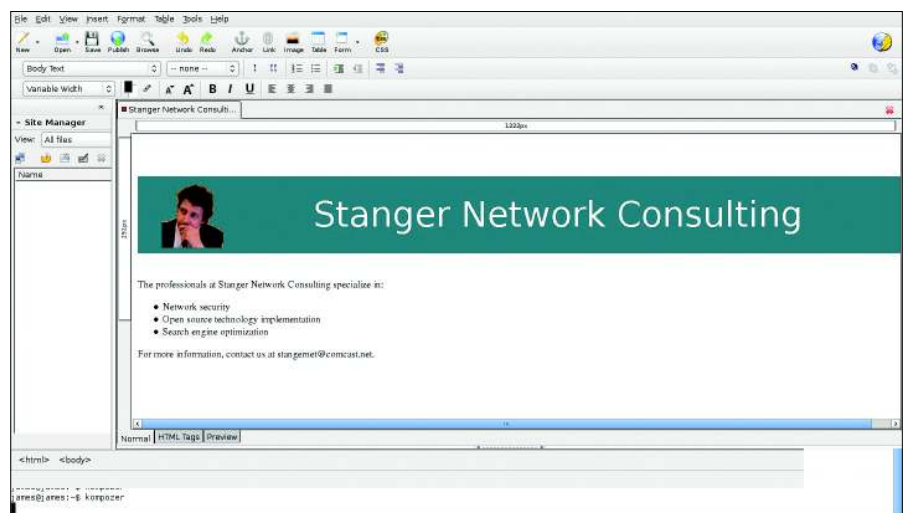


Figure 8: A web page in Kompozer after inserting an image.

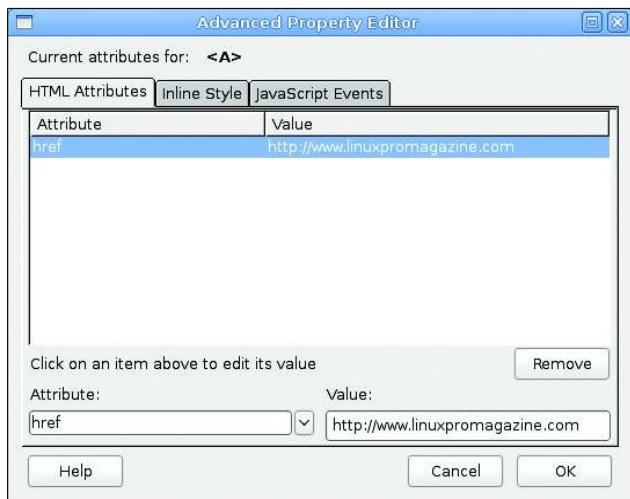


Figure 9: Inserting a link in Kompozer.

window, as shown in Figure 8. If the image doesn't appear, try viewing the code in a browser by clicking on the *Browse* button.

Hyperlinks

Hyperlinks are the live links familiar to all web users. If you click on a link, the browser jumps to another web page or marker. To insert a hyperlink in Kompozer, first highlight the text you want to turn into a link, then either click on the *Link* button or go to *Insert | Link* to bring up the *Link Properties* window. Next, click on the *Advanced Edit* button to bring up the Advanced Property Editor window, shown in Figure 9.

In the *Attribute* field in the lower-left portion of the window, use the drop-down box to select *href* (the attribute associated with a hyperlink). Then, in the *Value* field, enter the URL you want to link to.

To return to the page to see that the text you've highlighted is now a hyperlink, click the *OK* button twice.

HTML Validation

After you finish creating your web page, you're still not ready to post it. First you need to prove that the code you generated is compliant with the World Wide Web Consortium (W3C) standards. In addition to the

troubleshooting benefits of validation, validated pages render better in more browsers, and search engines tend to give validated pages a higher rank.

Various applications and services are available to help you get your code to pass W3C muster, including The W3C Markup Validation Service [11], Weblint [12], and Firefox add-ons such as Relaxed [13] and Total Validator [14].

To validate your web page with the W3C Markup Validation page, point your web browser to <http://validator.w3.org>.

The Validation page gives you the choice of entering a URL to a website or uploading a file. Because you haven't published this page yet, I'll assume that you're going to upload a file. As shown in Figure 10, click on the *Validate by File Upload* tab, then click on the *Browse* button.

To check a file, highlight it, click *Open* to return to the W3C page, then select *Check*. If your page was created with proper HTML or XHTML code, you will see a screen similar to that shown in Figure 11.

If your code has any compliance issues (and often it does), you will see a list of errors. Troubleshooting these errors takes quite a bit of time because the W3C service is not particularly intuitive with its hints and directions. Nevertheless, validation is an important step in ensuring that your web page will load properly in a variety of browsers.

But don't be happy just because your pages validate: It is also important to check the pages directly by opening them in a variety of web browsers.

Getting to the Server

Once you have created and tested your web page, it is time to post it on the staging server. In most cases, you won't be building your own whole new server but instead will be uploading files to an existing server system. Many Internet providers offer a small amount of web server space for subscribers, and most companies have access to servers that support a permanent web presence.

If you want to use the Apache web server included with Ubuntu to test and experiment with your HTML files, issue the following command to start Apache:

```
sudo /etc/init.d/apache2 start
```

The `/var/www/` directory is the default location for all files on the web server. All you have to do is place any file you want the server to render in this direc-

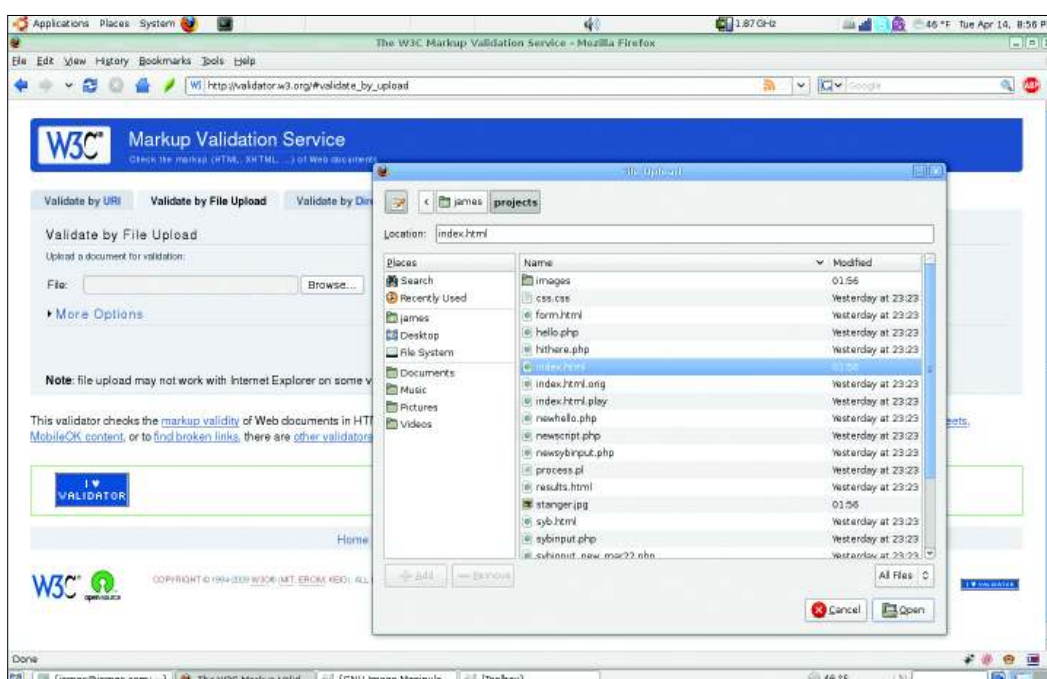


Figure 10: Choosing a file for validation.

tory. To place files in this directory, you need root privileges.

The default document is `index.html`. The server will automatically present this file to anyone who points their web browser at the server instance, so web developers typically give this name to the file that will serve as the homepage.

One popular way to copy the web files to the server is by entering the `cp` command in a terminal window:

```
sudo cp -r /var/www/projects/* /var/www
```

Because the root account owns the `/var/www/` directory, you need to use the `sudo` command.

If you need to send the files across the Internet, you can use ftp or even email to deliver the pages to the server.

Most HTML editors have a feature that lets you publish your site directly to the web server from within the editor. This feature is basically nothing more than an ftp client that you can configure to publish pages with one click. In Kompozer, for example, all you have to do is hit the *Publish* button on the taskbar or go to *File | Publish*. When you enter the neces-

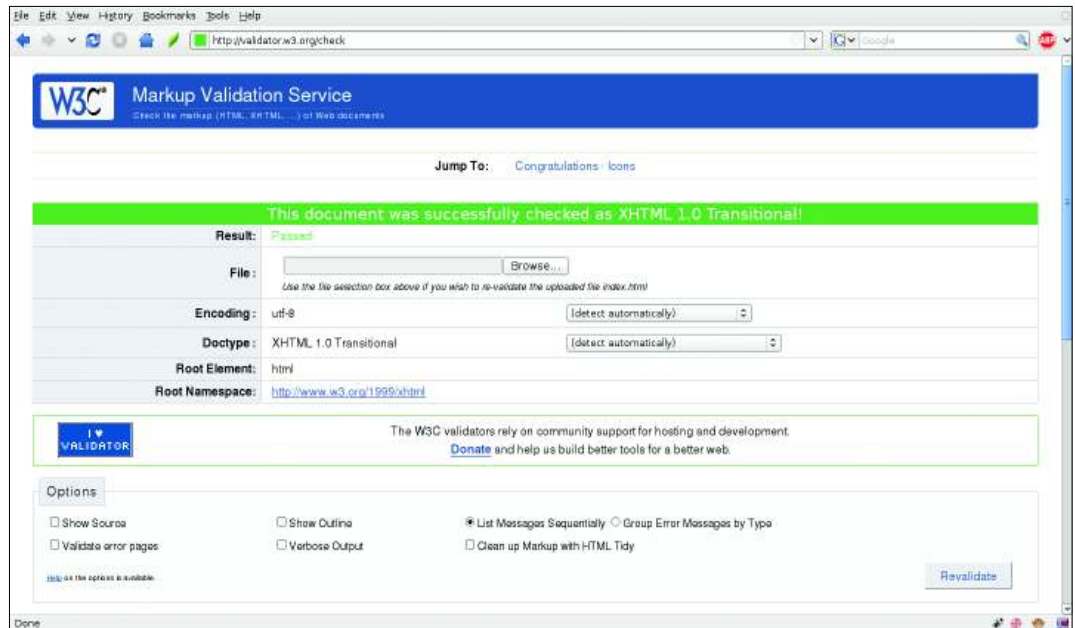


Figure 11: HTML code that has passed validation.

sary information in the Publish Page window, you are on your way to publishing the site.

Conclusion

Ubuntu provides a powerful, free development platform for building websites. In today's economy, the tools described in this article are more than enough for even a seasoned developer to create compelling, industry-standard web pages.

The next step for you to take is to finish establishing your development environment. Experiment with the tools that I've introduced here. Then, take a look at the myriad sites and education offer-

ings out there about web design. Sites such as *ciwcertified.com*, *how-to-build-websites.com*, and others can provide you with excellent ways to further your skills. ■

Creating a Simple Header Bar

GIMP is the de facto open source standard for creating standard RGB (Red Green Blue) files. GIMP is not exactly Adobe Photoshop, but it's just about good enough. GIMP supports several image file formats, including JPEG, GIF, and PNG.

The steps I used for creating the simple header bar shown in Figure 8 are as follows:

1. In the GIMP File menu, choose *New* and then specify the pixel dimensions. The size for this header is 1257x124 pixels.
2. Click on *Advanced Options* and then specify the X and Y resolution. Here, I will use 72 by 72 dots per inch (dpi) for the header, which is not ideal for images containing faces but is suitable for basic images used for navigation or background. If you specify a higher

resolution, the image might be clearer, but the page will take more time to download and render. Even in these days of increased bandwidth, a slight increase in download time could still make users impatient.

3. To specify a background color, choose the "bucket fill" tool and a color.
4. For the header, select the text feature and a color. Here, I will use white.
5. A second, already-existing image can be added to the image you are creating. For this example, I will drag a transparent GIF photo image onto the bar. "Transparent" means that the background of this new image will be the same color as the image behind it.
6. Now save the image as a JPEG by simply adding the `.jpg` ending to the file name.

INFO

- [1] Firebug: <https://addons.mozilla.org/en-US/firefox/addon/1843>
- [2] Web Developer: <https://addons.mozilla.org/en-US/firefox/addon/60>
- [3] YSlow: <https://addons.mozilla.org/en-US/firefox/addon/5369>
- [4] SenSEO: <https://addons.mozilla.org/en-US/firefox/addon/9403>
- [5] Bluefish: <http://bluefish.openoffice.nl>
- [6] Quanta Plus: <http://quanta.kdwebdev.org/>
- [7] Kompozer: <http://www.kompozer.net>
- [8] Open Source Web Design: <http://www.oswd.org>
- [9] Open Source Templates: <http://www.opensourcetemplates.org>
- [10] Kompozer issues: <http://ubuntuforums.org/showthread.php?t=1094500>
- [11] The W3C Markup Validation Service: <http://validator.w3.org>
- [12] Weblint: <http://www.w3.org/Tools/weblint.html>
- [13] Relaxed: <https://addons.mozilla.org/en-US/firefox/addon/3939>
- [14] Total Validator: <https://addons.mozilla.org/en-US/firefox/search?q=html+validator&cat=all>



Rolling your own Debian packages

GIFT WRAPPED

Just a couple of steps are all it takes to convert a source code archive into a full-fledged Debian package.

BY CARSTEN SCHNOBER

The compiler was once regarded as an exotic tool for hard-core developers, but building programs is now a common task for average Linux users. The three classic commands *configure*, *make*, and *sudo make install* will help you convert the source code downloaded from a promising project website into a binary application ready to execute on your own system.

The problem is, if you build a program yourself, you won't be able to manage, track, and update the application through Ubuntu's handy package management system.

If you can't find an Ubuntu package for the program you want to install, but you still want the chance to manage the program using Ubuntu's package management tools, you can always convert the program into the Debian package format used with Ubuntu. In this article, I'll walk you through the steps required

to convert a program into the Debian package format.

CheckInstall

The CheckInstall [1] program (Figure 1) offers the easiest approach to creating the DEB packages used with Ubuntu, as well as RPM packages used with other Linux variants. CheckInstall is easy to use, although it tends to fail with packages that extend the operating system functionality by adding kernel modules. If you prefer DEB packages that will run on systems other than Debian, you need more than the basic packages created by CheckInstall. Debian-based distros offer a collection of tools, mainly found in the *dpkg-dev* package, that facilitate the task of packaging an application.

Original Beats Imitations

Debian packages are quite easy to build as long as the source code directory is

ready. All the required information must be located in the *debian* subdirectory of the source code directory. If the distribution has a prebuilt DEB package with the software you want to compile, the information is probably already available. To recompile this version – for example, with support for additional features – first download the matching source code using *apt-get source Packagename* (without *sudo*). You might need to enable the *Source code* package source in Synaptic.

After you download the files to a directory, you will need to change to this directory to run the command. You should have three new files and a directory. The filenames comprise the software name, the package version number, and possibly an internal revision number, such as *ubuntu2*. One of the three files is the original archive with the *.orig.tar.gz* suffix. The file with the *.diff.gz* suffix contains the changes the

distributor has made to the original source code to adapt the package to match the system. Then there is a `.dsc` file containing a description with details of the supported processor architectures and the package maintainer. The directory created by `apt-get source` also contains the `debian` directory.

In the simplest case, you could now just make any changes you need; for example, adding a program library to the software to extend its functionality. You can add more features to the source code or install patches with changes. To create a new Debian package, issue the `dpkg-buildpackage` command in the source code directory (the command is part of the `dpkg-dev` package).

To build a package using this approach, you actually need administrative privileges; however, if you install the `fakeroor` package, `dpkg-buildpackage` will automatically use it if you try to run the program without administrative privileges. `Fakeroor` lets you build Debian packages with write privileges for the

source code directory and its parent directory.

First, the program checks to see whether the tools and libraries it needs for building the package exist on the system. If not, it issues a list of missing packages you need to install. Otherwise, the script will call `configure` and then `make` before finally creating a DEB package that it stores in the parent directory. To install the package, you need to launch the package manager; at the command line, the following command handles this:

```
$ sudo dpkg -i package_name.deb
```

Documenting the details is a good idea. The `changelog` file in the `debian` directory is the place to document your work,

```
File Edit View Terminal Help
root@system:~/liferea/liferea-1.4.26# sudo checkinstall
checkinstall 1.6.1. Copyright 2002 Felipe Eduardo Sanchez Diaz Duran
This software is released under the GNU GPL...

Please write a description for the package.
End your description with an empty line or EOF.
>> A simple and free Gtk desktop news aggregator for Unix and Linux
>>
*****
**** Debian package creation selected ****
*****

This package will be built according to these values:

0 - Maintainer: [ root@system ]
1 - Summary: [ A simple and free Gtk desktop news aggregator for Unix and Linux ]
2 - Name: [ liferea ]
3 - Version: [ 1.4.26 ]
4 - Release: [ 1 ]
5 - License: [ GPL ]
6 - Group: [ checkinstall ]
7 - Architecture: [ amd64 ]
8 - Source location: [ liferea:1.4.26 ]
9 - Alternate source location: [ ]
10 - Requires: [ ]
11 - Provides: [ liferea ]

Enter a number to change any of them or press ENTER to continue:
```

Figure 1: CheckInstall collects information about the program, but it does not define any dependencies.

but make sure you stick to the required format; otherwise, your package-building activities are doomed to failure.

Updates

In real life, program packages included by distributors often do not come with

Table 1: Files in the debian Directory

Filename	Filename and Path after Installation	Function
<code>README.Debian</code>	<code>/usr/share/doc/package_name/README.Debian</code>	Documents differences between the original source code and the “debianized” version
<code>changelog</code>	<code>/usr/share/doc/package_name/changelog.Debian.gz</code>	Logs the changes made to the package
<code>copyright</code>	<code>/usr/share/doc/package_name/changelog.Debian.gz</code>	License under which the software in the package is released
<code>control</code>	–	Information for the package: name, version, dependencies, description, etc.
<code>rules</code>	–	Rules for creating the package
<code>conffiles.ex</code>	–	Configuration files contained in the package
<code>cron.d.ex</code>	<code>/etc/cron.d/package_name</code>	Actions and programs that the program tells cron to execute regularly (in cron format)
<code>dirs</code>	–	List of directories required by the program, without leading slashes; as in <code>usr/bin</code>
<code>docs</code>	–	List of files to be copied to the <code>/usr/share/doc/package_name/directory</code>
<code>emacsen-install.ex</code>	–	Emacs-specific (optional)
<code>emacsen-remove.ex</code>	–	Emacs-specific (optional)
<code>emacsen-startup.ex</code>	–	Emacs-specific (optional)
<code>init.d.ex</code>	<code>/etc/init.d/package_name</code>	The init script for starting and ending any system services contained in the package
<code>manpage.1.ex</code>	<code>/usr/share/man/man1/program_name.1.gz</code>	Man page for the program
<code>menu.ex</code>	<code>/usr/share/menu/package_name</code>	Information for a program entry in the start menu
<code>watch.ex</code>	–	Configuration file to update the package using <code>uscan</code> and <code>uupdate</code>
<code>preinst.ex</code>	–	Contains the commands the package manager executes before installing the package
<code>postinst.ex</code>	–	Contains the commands the package manager executes after installing the package
<code>prerm.ex</code>	–	Contains the commands the package manager executes before removing the package
<code>postrm.ex</code>	–	Contains the commands the package manager executes after removing the package



Figure 2: The update command unpacks a source code archive and creates a debian subdirectory based on the content of the existing package.




Figure 3: The dh_make command Debianizes a source code directory, creating a debian subdirectory where it creates and generically populates files.

the latest version; you need a new DEB package based on the source code for the updated program version. The `uupdate` command from the `Devscripts` package helps you update the source code.

To use `uupdate`, download the tarball with the new archive, but do not unpack it. Then run `apt-get source program` (again without `sudo`) to download your distributor's version of the same package. This will give you both a complete `debian` directory with information you can recycle, as well as (typically) the changes made by the distributor to modify the system for your distribution.

Now change to the `program version` directory created by `apt-get source` and run the `uupdate -u` command, adding the filename of the archive containing the new version of the program (Figure 2):

```
$ uupdate -u Z
path/program-NewVersion.tar.gz
```

The command reads the changes made by the distributor from the `program.diff`

.gz file and attempts to add them to the new version. This approach will only work if the modified files are not much different from their counterparts in the previous version. If just a couple of the patches fail, the update process will carry on without them. `uupdate` also transfers information, such as the package description and dependencies, from the original package to the new version.

The call to `uupdate` unpacks the tarball with the updated program version, which contains a new source code directory, including a `debian` subdirectory with the package information gleaned from the original package. Next, run `dpkg-buildpackage` in the new source code directory, as described previously. This step starts the build, which gives you a full-fledged DEB package. Again, the package management tool is required to install the package:

```
dpkg -i package_name.deb
```

New Buildings

If you do not have a previous version of the program to build on, your only option is a completely new DEB package. To start, unpack the source code archive, which will not contain any distribution-specific information. To "Debianize" the source code, you must first install the `dh-make` package. Then run `dh_make` with the `-f` parameter in the source code directory, passing in the path to the original source code archive. So that other users can contact you, you also need to enter an email address for the package, which is defined with the `-e` parameter:

```
$ dh_make -e mail@address.com -f
../program-version.tar.gz
```

Next, you must decide what type of package you are building (Figure 3). In the easiest case, you enter `s` for a *single binary* package. The other options are `m` for *multiple binary*, `l` for *library*, and `k` for *kernel module*.

The CDBS (`b`, Common Debian Build System, [2][3]) option is useful mainly for more complex packages that you will want to use and update over an extended period of time; however, this option does require a deeper understanding of the DEB package system.

debian Directory

After you run `dh_make`, take a look at the files in the `debian` directory (Table 1). The files created here contain the placeholders and generic information you will need to modify to match the package.

The `control` file plays a fundamental role because it contains all the critical package information. You can enter the software category in the `Section:` line; Table 2 shows your options.

In many cases, an additional `Priority:` line is useful. The `Priority:` line describes how important the software is for the

Table 2: Software Categories for "Section:"

Designation	Purpose
<code>admin</code>	System management programs
<code>base</code>	Basic packages
<code>devel</code>	Packages for software developers
<code>doc</code>	Documentation
<code>libs</code>	Program libraries
<code>mail</code>	Email software
<code>net</code>	Network software
<code>x11</code>	Other graphical programs

INFO

- [1] CheckInstall: <http://asic-linux.com.mx/~izto/checkinstall>
- [2] CDBS: <http://build-common.alioth.debian.org>
- [3] Debian packages via CDBS: <http://www.ngolde.de/cdb.html>

system; in the case of packages you build yourself, the priority will typically be *optional* or *extra*. The dependencies are specified in the *Depends:* line, where *dh_make* has already created an entry:

```
 ${shlibs:Depends}, 2
 ${misc:Depends}
```

These variables save you some work; when you compile, the *dh_shlibdeps* program automatically replaces the variables with the required packages.

If you want to add recommended packages, insert a *Suggests:* line into the *control* file;

```
Suggests: libqt0-ruby1.8
```

The packaged program will not strictly require the programs and libraries stated in the *Suggests:* line, but in many cases, the recommended packages will enable features that users would otherwise have to do without.

All that remains is to fill out the program *Description:* field. The first line should contain a short description of no more than 60 characters; below it, there is enough space for more details. Indent each line with a blank so that the package builder will know that the lines belong to the *Description:* field.

Home Straight

The *rules* file is another important inhabitant of the *debian* directory. The *dpkg-buildpackage* command, which finally triggers the package build, takes its instructions from the *rules* file. For simple programs, the rules generated by *dh_make* will work without modification. If you are packaging more complex software or program libraries, you can add parameters, such as compiler options, to the *rules* file.

The *copyright* file contains the names and email addresses of the programmers and copyright owners, and the license under which the software package is published. Indent the text with four blanks. If you are using GPL licensing, you don't need to add the whole text, since it is available on any Debian-based system in the */usr/share/common-licenses/GPL* directory. In this case, you can simply add a reference to the license text in *copyright*.


Checking the *changelog* file for any new package is a good idea. The first time you build a package, you normally won't need to change the log and can just delete the file. In all other cases, you will need to edit the file. The *README.Debian* file is where you add instructions that relate to the DEB package in general and not to the software it contains.

Small Steps

If you build software yourself, the small effort of building matching packages will give you enormous convenience benefits. The *uupdate* utility updates on the basis of an existing package, and *dh_make* creates the *debian* directory. With just a little practice, a final call to *dpkg-buildpackage* quickly becomes routine.

Seeing your first home-grown package install is a really great feeling! ■

License Free Virtualisation



At Digital Networks, we build, test, sell and support virtualisation platforms. Our platforms focus on technical solutions and are designed to reduce your costs and simplify management.

We've been building Linux systems since 1998 and we offer standards-based platforms that integrate with Linux, Windows, Unix and Mac networks.

We offer KVM, OpenVZ, Xen and VirtualBox technologies. We focus on license free virtualisation using open source software. We can provide load balanced high availability solutions without the huge licensing costs of proprietary virtualisation platforms.

For more information, visit www.dnuk.com

Premium Ubuntu PCs

We also sell high-end Ubuntu PCs. Visit our site to have a look at our new AMD Phenom II X4 and Intel Core i7 machines.

These are pre-installed with Ubuntu 9.04.



 Digital Networks
United Kingdom

www.dnuk.com sales@dnuk.com 0161 366 6600



Mike Basinger is an administrator on the Ubuntu Forums, where he helps users with Ubuntu questions every day. He works for the Ubuntu community by serving on the Ubuntu Community and Forum Councils. Mike works at the University of Utah Marriott Library in Salt Lake City, Utah, as an IT manager. Send your Ubuntu questions to: QandA@ubuntu-user.com

Q&A with Ubuntu Forums admin Mike Basinger

ANSWER MAN

```
sudo apt-get install \
evolution-mapi
```

This will only work if your Exchange 2007 server has MAPI enabled on it.

Installing from USB

? I have an Asus Eee PC 901 and want to install Ubuntu on it, but it has no CD/DVD drive. How can I copy the Ubuntu Live CD ISO [image] to a USB flash drive?

💡 The ability to create a bootable Ubuntu USB installer has been included on the Ubuntu Live CD since the release of Ubuntu 8.10. All you need to make a USB installer is a USB flash drive with at least 700MB of free space, an Ubuntu 8.10 or newer Live CD, and a computer that you can boot into the Ubuntu Live CD.

Place the Ubuntu Live CD in your computer CD/DVD drive and boot from the Live CD. At the prompt, choose *Try Ubuntu without installing*. Once the Ubuntu Live session boots, you should see an Ubuntu basic desktop.

Now plug your flash drive into your computer and go to the top menubar and choose *System | Administration | Create a USB startup disk*. Next, select your Live CD in the top selection box



Figure 1: Creating a USB startup disk.

and your flash drive in the bottom selection box and click on *Make Startup Disk* (Figure 1). You can also run this same program by typing *usb-creator* at the command line on any existing Ubuntu 8.10 or newer machine.

Fixing Ubuntu Boot Menu

? I still dual boot my machine between Windows XP and Ubuntu. I recently had to reinstall Windows XP on my computer. After reinstalling Windows, my Ubuntu boot menu was gone, and I could no longer boot into Ubuntu. How can I recover my Ubuntu boot menu so I can dual boot between Ubuntu and Windows XP again?

💡 When you reinstalled Windows, the Ubuntu bootloader, *grub* got overwritten. Fixing this is easy – you just need an Ubuntu install CD or Live CD. Once the Ubuntu Live session boots, you should see an Ubuntu basic desktop.

Now open a terminal window, *Applications | Accessories | Terminal* and type the following:

```
sudo grub
```

This will open the grub command-line interface. At the *grub* prompt, enter the following commands:

```
find /boot/grub/stage1
```

This will tell you what hard drive and partition grub is on. For this example, I will use *(hd0,0)*, the first partition on the first hard drive, which is a very common setup.

Now you want to repair the grub bootloader:

```
root (hd0,0)
setup (hd0)
exit
```

The *root* command points to the proper area of the hard drive, and the *setup*

Corporate Email

? My company uses Microsoft Exchange server for email. Does Ubuntu have any mail clients that work with an Exchange server?

💡 If you don't need calendaring, you can use almost any IMAP (Internet Mail Access Protocol) mail client, such as Thunderbird, Mutt, Evolution, or KMail, to connect to your mailbox. Ask your admin whether your Exchange server is set up to work with IMAP. If you do need calendaring, you can try using the Exchange connectors included with the Evolution mail client. Ubuntu automatically installs the Exchange connector needed for Microsoft Exchange 2003 or earlier. If your workplace has a Exchange 2007 server, you need to install the MAPI (Messaging Application Programming Interface) plugin for Evolution, which is a new feature in Evolution 2.26:

command installs grub to the MBR (Master Boot Record).

To get your Ubuntu boot menu back, close down your Ubuntu Live CD session, remove the CD, and reboot your computer normally.

Playing DVDs



I would like to play DVDs on my Ubuntu laptop when I'm on a flight, but it does not work when I insert a DVD. Is there anyway to make DVD playback work in Ubuntu?



Most store-bought DVDs are encrypted with CSS (Content Scrambling System), which attempts to restrict the software that can play a DVD. The ability to decrypt DVDs is not included by default on Ubuntu because of legal restrictions in some countries. To enable DVD playback (assuming it is legal in your area and you have a legal right to access the recording), you need to install the `libdvdcss` library. From Synaptic or a terminal, install `libdvdread4/` to get `libdvdcss`:

```
sudo apt-get install libdvdread4
```

Then from a terminal, execute the following command to install `libdvdcss`:

```
sudo /usr/share/doc/2
libdvdread4/install-css.sh
```

Games and Compiz



I really like the visual effects in Compiz, but it causes slow-downs and screen corruption when I play games in Ubuntu. Is there a quick and easy way to turn off Compiz temporarily when I play games on my computer?



Compiz can be turned off with the Appearance control panel, *System | Preferences | Appearance*, by clicking on the *Visual Effects* tab and selecting *None*. The problem with this method is that any special settings that were changed in Compiz will be lost when you re-enable it.

A better way to turn Compiz off and on is to use the `fusion-icon` program. The `fusion-icon` program sits in the notification area in your Gnome panel and controls Compiz. Install `fusion-icon` with Synaptic or `apt-get`:

```
sudo apt-get install fusion-icon
```

Now you can launch `fusion-icon` at login by creating a new startup program in *System | Preferences | Startup Applications*. Be sure to use `fusion-icon -n` so that it does not change the window manager when it starts.

NVIDIA Graphics Card



I bought a new NVIDIA card and it is not supported by the NVIDIA drivers that come with Ubuntu. How do I install the latest driver from the NVIDIA website on Ubuntu?



Installing NVIDIA drivers manually is a little different on Ubuntu than on other Linux distributions. First you need to download the latest NVIDIA driver from their website that supports your card. To compile the driver, you need some extra software installed in Ubuntu:

```
sudo apt-get install 2
build-essential linux-headers-2
`uname -r`
```

Then you need to install the 3D NVIDIA driver that Ubuntu installs:

```
sudo apt-get remove 2
nvidia-glx nvidia-glx-legacy 2
nvidia-glx-new nvidia-settings
```

The Ubuntu Linux kernel includes the 2D NVIDIA driver, which must be disabled at boot-up. Just edit the `/etc/default/linux-restricted-modules-common` file:

```
gksudo gedit 2
/etc/default/linux-restricted-2
modules-common
```

and change the line `DISABLED_MODULES = "` to `"DISABLED_MODULES = "nv nvidia_new"`.

Be sure to back up your `xorg.conf` file and exit the X Window System completely.

```
sudo cp /etc/X11/xorg.conf 2
/etc/X11/xorg.conf.backup
```

For Ubuntu, type

```
sudo /etc/init.d/gdm stop
```

and for Kubuntu, enter:

```
sudo /etc/init.d/kdm stop
```

Now log in to your machine at the terminal window and go to the directory in which you downloaded the NVIDIA driver. The `chmod` command will make the file executable:

```
sudo chmod +x NVIDIA*
```

Then, run the driver installer, follow the instructions, and reboot your computer after you finish.

If all goes well, your computer reloads X Window with full 3D support:

```
sudo sh NVIDIA*
```

Bug Out



I'm having problems with Firefox. It does not behave as it should and crashes on me occasionally. How do I report problems back to Ubuntu?



If you are having a problem with Ubuntu that you think the developers should know about, submit a bug report to Launchpad (<http://launchpad.net>). Launchpad provides the bug tracker, among other services, used by the Ubuntu community. Before you submit a bug to Launchpad, you will need to create an account at <https://launchpad.net/distros/ubuntu/+bugs/+login>. After you have created your account, you can use two automatic procedures to create your bug report. In most programs included in Ubuntu, go to *Help | Report a Problem*. This will automatically create a bug report with the necessary information that Ubuntu developers will need.

All you need to do is fill in a description of what is happening. At the command prompt, you can do the same thing with the `ubuntu-bug` command:

```
ubuntu-bug firefox
```

Before submitting a bug report, I would suggest searching Google, the Ubuntu Forums, and Launchpad to see whether this is a known issue. Hunting down a previous answer will help keep the developers from being overwhelmed by duplicate bug reports. ■

Launching programs and actions with Gnome Do

CENTRAL LAUNCHER

Redrex, Fotolia

Information at your fingertips with actions to match – Gnome Do merges search engine and application launcher functionality to create a new and convenient tool. **BY CHRISTOPH LANGNER**

Documents, applications, messages, and bookmarks tend to spread all over a system. Gnome Do [1] gives users a new launcher for Gnome and other desktops that not only finds information but also triggers actions on the basis of the results.

The flexible Gnome Do [1] application launcher is similar to the Quicksilver [2] launch tool for Mac OS X. Users press a keyboard shortcut to pop up a window in which they enter the name of the program they need.

Because Gnome Do auto-completes the entry, entering just part of the name is sufficient (Figure 1). For example, typing *text* will launch your text editor. Gnome Do bases its auto-complete suggestions on start menu entries.

Gnome Do can do more than just launch programs – it will find music in Rhythmbox (see Figure 2), Amarok, or Banshee databases or contacts in Evolution, Gmail, or the Pidgin instant messenger. If you want, you can tell Gnome Do to set up a chat session with the con-



Figure 1: Gnome Do auto-completes the entry. (Screenshots courtesy of David Siegel)

tact in Pidgin. The tool also supports file operations and much more. Various plugins give users a convenient approach to extending functionality.

Installation

Gnome Do is easy to install because Ubuntu includes all the necessary packages [3]. Simply run *Applications | Add/Remove...* and enter *gnome-do* into the *Search* box. Then mark the resulting entry *GNOME Do* for installation and click on *Apply Changes*.

After you finish the install, *Applications | Accessories | GNOME Do* takes



Figure 2: Gnome Do even finds your music.



Figure 3: Plugins extend Gnome Do's functionality.

you to a program launcher. First you will see the application window. If you click on the desktop or perform some other action, Gnome Do disappears into the background.

To move the tool to the foreground, you just press its configured keyboard shortcut – by default, that's Windows + Spacebar.

Session

After installation, Ubuntu automatically launches Gnome Do in the background at every start of the desktop. With the session management feature, you can undo and later reactivate this behavior if you want.

To remove Gnome Do from the list of programs that are started automatically, press *System | Preferences | Startup Applications* and search for the *GNOME Do* entry. Now you want to deactivate the checkbox next to it and click on *Close*. In the same way, you can later configure the desktop to run Gnome Do at startup again.

Gnome Do does not use a traditional application window. Assuming you have desktop effects enabled, Gnome Do does without a frame and the window is transparent.

To enable the desktop effects, click on the *Appearance* entry in the *System | Preferences* menu. Now, in

the *Visual Effects* tab, you can choose *Normal* or *Extra*.

Triple Functionality

The Gnome Do main window has three sections, two of which are visible on launching the tool. When you enter a search key, an action to match the search automatically appears in the second section. For example, if you enter *text*, Gnome Do automatically selects *Run* to execute a program. If you then confirm the action by pressing Enter, the Gnome text editor launches.

Because Gnome Do uses the names in the start menu to identify applications, it can cause some confusion. The Gnome text editor binary is called *gedit*. If you enter the name in Gnome Do, the program fails to find a match. To be on the safe side, stick to the names in the application menus.

If you enter the name of a file or directory in your home directory and press Tab to move to the action area, you can then press the Down arrow key to select other actions, such as moving or copying. After selecting an action, a third section appears and you can choose a target for the action.

These controls are typical of Gnome Do: You start by typing a search key and then press Tab to move on to the next section, where you press the Down arrow key to select an action.

Plugins

The recent Plugin improvements have made Gnome Do a flexibly extendable tool [4]. Each plugin has a library with a

file extension of *.dll*, which is fairly surprising for Linux users. Simply drop the plugin into the *.local/share/gnome-do/plugins* directory below your home directory and relaunch the program to allow Gnome Do to detect the plugin. This gives users a simple approach to adding a variety of functions.

Gnome Do 0.5 introduced a new preferences window and plugin manager (Figure 3), which helps you browse, download, install, and enable or disable plugins from within the tool. Additional plugin information is available on a wiki page, and you can do plugin-specific configuration.

Gnome Do will search Firefox or Epiphany bookmarks and open links in your browser, and it will also search for music in your Rhythmbox, Amarok, or Banshee database and add this to a playlist (see Figure 4). Also, you can add Pidgin or Evolution contacts to the launcher.

Community Contributions

New plugins are released almost daily. Gnome Do provides a community plugin category with plugins from a variety of contributors. New community contributions include Skype, Twitter, Flickr, and Google Calendar plugins, as well as the WindowManager plugin that lets you arrange windows on your desktop. Developers who enjoy working with Mono can also refer to the detailed documentation to implement their own ideas [5]. ■



Figure 4: Gnome Do searches the Rhythmbox music database for artists and albums and offers actions relevant to the current context.

INFO

- [1] Gnome Do: <http://do.davebsd.com/>
- [2] Quicksilver: http://docs.blacktree.com/quicksilver/what_is_quicksilver
- [3] Installation: http://do.davebsd.com/wiki/index.php?title=Installing_Do
- [4] Gnome Do plugins: <http://do.davebsd.com/wiki/index.php?title=Category:Plugins>
- [5] Developing plugins: http://do.davebsd.com/wiki/index.php?title=Getting_started_with_writing_plugins

AUTHOR

Christoph Langner works in test management for PTV in Karlsruhe, Germany, and has played an active role in the open source software community for many years.

Getting the most out of FreeMind

MIND TRICKS

Learn how to master the advanced features of the FreeMind mind-mapping application to get the most out of this useful tool. **BY DMITRI POPOV**

Although you might think that FreeMind [1] is just a mind-mapping application designed to help creative professionals organize their thoughts and ideas, in reality, it sports so many nifty features that it could prove to be an indispensable productivity tool no matter what you do for a living.

Because FreeMind is a rather straightforward application, you can learn how to use its basic features in a couple of hours. But if you want to get the most out of this excellent application, you might want to master its more advanced features.

Basics

The Ubuntu 9.04 package manager only contains FreeMind version 0.7. To install it, use the *Applications | Add/Remove...* dialog.

On the FreeMind homepage you will find the most recent program version, currently 0.8.1, which brings a number of new features. Just click on Debian

Linux and download the file *freemind_0.8.1-2_all.deb*. A double-click on it starts the installer for external packages. The plugin packages provide additional features but are not necessary.

The package installer automatically takes care of the additional necessary files. The FreeMind developers recommend Sun Java instead of the free GCJ. To get it, install the *sun-java6-jre* package with the Synaptic package manager or run the *sudo apt-get install sun-java6-jre* command in the terminal. Once the package is installed, use the *sudo update-alternatives --config java* command to set Sun JRE as the default environment on your system.

If you don't mind living on the bleeding edge, you might want to opt for the latest test release of FreeMind that comes with a few nifty features not found in the stable version of the application. No official distribution packages are ready yet, so you have to download the zip archive containing the latest de-

velopment version of FreeMind (currently, it's 0.9.0) and unpack it.

In the terminal, switch to the *freemind* directory and run the *sh freemind.sh* command to launch the application.

Using Shortcuts

Because it is a graphical application, FreeMind lets you work with mind maps by the point-and-click method. It also supports keyboard shortcuts, which allow you to perform virtually any operation without touching the mouse.

The main advantage of keyboard shortcuts is speed: Hitting a key combination is much faster than finding and moving the mouse. It might take you some time to learn FreeMind's shortcuts, but once you've mastered them, you can create and navigate mind maps with consummate ease and speed.

These keyboard skills can come in particularly handy when you are brainstorming a problem or taking notes during a lecture and you don't have time to

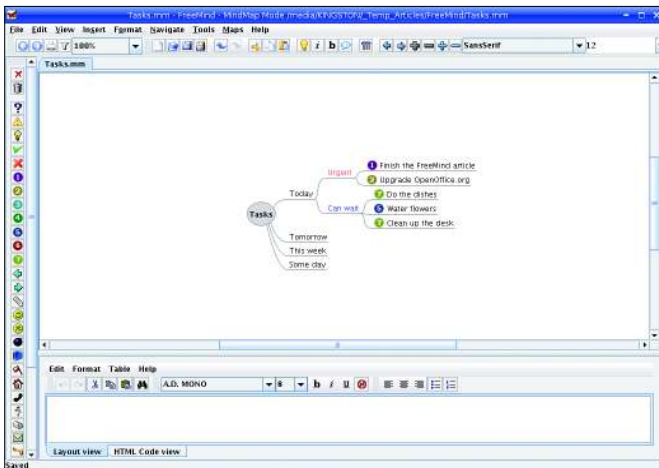


Figure 1: Simple task mind map.

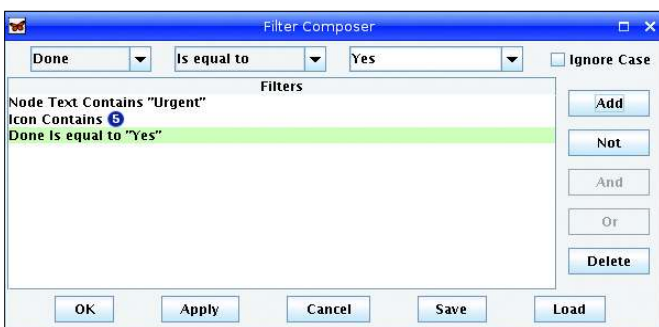


Figure 2: Defining filters.

fumble with the mouse or the trackpad on your laptop. The full list of FreeMind keyboard shortcuts is rather long, but you can start by learning the most essential ones:

- Enter: Add sibling node
- Insert: Add child node
- F2: Edit selected node
- Alt + Enter: Edit selected node with the rich text editor
- Space: Fold/unfold the selected node
- Ctrl + Up arrow: Move node up
- Ctrl + Down arrow: Move node down
- Ctrl + Shift + K: Add link with the file chooser
- Ctrl + K: Add link by entering it manually

Usage Scenarios

FreeMind is not only good for mapping your ideas, with a bit of creativity, you can use FreeMind for pretty much anything. For example, FreeMind has all the features you need to create a task manager. To start, create a mind map that represents the overall structure of a task manager.

The exact structure depends on how you usually manage your tasks, but it could look like the example in Figure 1.

Because FreeMind supports drag-and-drop, you can move your tasks between categories with the mouse. For example, if you don't have time to do the dishes today, you can move the task to tomorrow by dragging it onto the *Tomorrow* node. Also, you can use the available icons to assign priorities to the tasks quickly; right-click on the task and choose the desired priority icon from the icons list.

The Filter feature can help you manage complex task lists (see Figure 2). Say you want to display only the tasks with Priority 1. Just press the *Edit* button (it looks like a pencil) in the Filter toolbar, and use the available dropdown lists to define the following rule: *Icon - Contains - 1*. Press the *Add* button to add the rule to the list, then press *Apply* to enable it.

In a similar manner, you can create a filter that shows only tasks in the Urgent category by creating the *Node - Contains - Urgent* rule; just remember to tick the *Show Descendants* checkbox to display the tasks in the Urgent category.

Attributes is another powerful feature you can use to keep tabs on your tasks. With it, you can add the *Done* attribute with values of *Yes* and *No*, which you can use to mark finished and unfinished tasks (Figure 3). Just choose *Tools | Attribute Manager*, and press the *Edit* button in the *All Attributes* row. Next, type the name of the attribute (i.e., *Done*) and press the *Add* button. Then press the *Edit* button next to the *Done* attribute and add the *Yes* and *No* values to it. To assign the defined attribute, select the desired task and press Alt + F9 (or right-click on the task and choose *Edit Attributes*). Add the *Done* attribute and set its value.

Thanks to the inclusion of a rich text editor that lets you add notes to the nodes, FreeMind makes a powerful and flexible note-taking tool. For example, if you are attending a conference, you can create a mind map containing separate nodes for each conference track, and you can specify different attributes that allow you to identify each track quickly (e.g., "Development," "Marketing," "Birds of a Feather," etc.). Then, to add richly formatted notes to each node, you can use the text editor at the bottom of the main window.

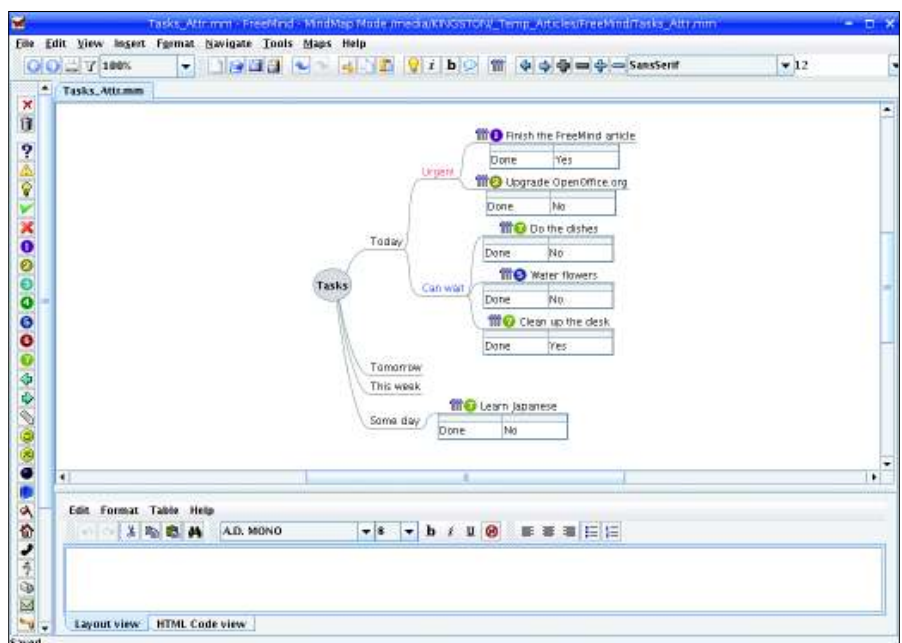


Figure 3: Using attributes to keep tabs on tasks.

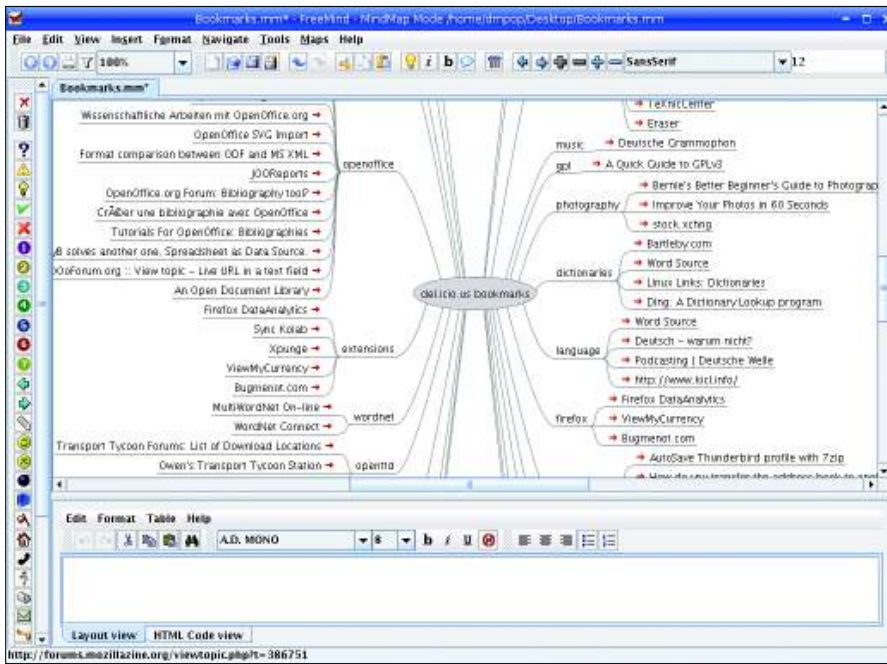


Figure 4: Displaying your Delicious bookmarks as a mind map.

FreeMind can also encrypt individual nodes and entire mind maps, making it suitable for storing private information. With this feature, you can turn a mind map into a password and account manager. Say you want to use a mind map to store your account information for different websites. Just choose *Tools | Insert Encrypted Node* and specify a password.

Once the encrypted node is added, you can add child nodes to it for each website account. Then, you can use the text editor to add notes containing login and other important information. Next time you open the mind map, you must enter the specified password to access the encrypted node and its children. The ability to encrypt individual nodes can come in handy when you want to create a mind map containing both public and private information.

If you prefer to keep all confidential data in a separate file, you can create an encrypted map with *File | Create Encrypted Map*.

With third-party tools, you can put FreeMind to some

less conventional uses, too. If the way Delicious displays bookmarks doesn't tickle your fancy, you can add a whole new dimension to your bookmark collection by turning it into a mind map with a clever utility called *Delicious-Mind* [2]. Just download the *deliciousmind.zip* archive and unpack it into your home directory.

Next, you have to export bookmarks from your Delicious account. To do this, point your browser to <http://del.icio.us/api/posts/recent?count=10000>, enter your username and password when prompted, then save the page as *bookmarks.xml* in the *deliciousmind* directory (in Firefox, choose *File | Save Page As*). In the terminal, switch to the *deliciousmind* directory and run the *sh deli-*

ciousmind.sh bookmarks.xml bookmarks.mm command. This converts the XML file into a mind map, which you can open in FreeMind (Figure 4).

Getting Physical

To provide a more efficient way of applying formatting to nodes in a mind map, FreeMind sports so-called physical styles, which you can use to specify styles and their properties, such as node color, type, font, borders, and so on. To do this, choose *Format | Physical Style | Manage Patterns*. This opens the Patterns dialog window (Figure 5), where you can tweak the existing styles. To apply a physical style to a node, right-click on it and select the desired style from the *Physical Style* list.

Exporting Mind Maps

FreeMind can export mind maps in a wide range of formats, including PDF, ODT, PNG, SVG, and interactive HTML pages, as well as Java and Flash applets, and it can even export your mind maps as TWiki pages and TaskJuggler tasks. But that's not all.

The *mm2odp* [3] and *mm2s5* [4] utilities, for example, let you turn FreeMind mind maps into OpenOffice.org Impress and S5 presentations, whereas the *mm2notes* [5] tool can transform a mind map into hierarchical notes. All three utilities are written in Python, and they are easy to install and use. ■

INFO

- [1] FreeMind: freemind.sourceforge.net/wiki/index.php/Main_Page
- [2] DeliciousMind: www.blainekendall.com/deliciousmind
- [3] *mm2odp*: code.google.com/p/mm2odp
- [4] *mm2s5*: code.google.com/p/mm2s5
- [5] *mm2notes*: code.google.com/p/mm2notes

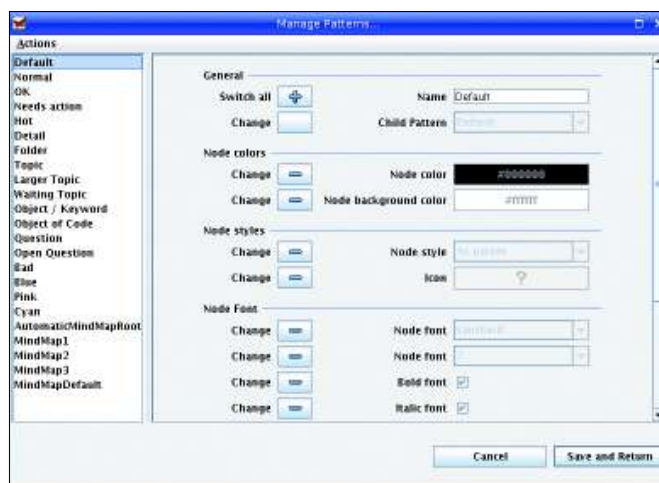


Figure 5: Working with physical styles.

THE AUTHOR

Dmitri Popov holds a degree in Russian language and computer linguistics. He has been working as a technical translator and freelance contributor for several years. He has published over 500 articles in Danish, British, US, and Russian magazines and websites.



Creative prose writing with Celtx

WRITERS UNBLOCKED



Sandor Jachai, Fotolia

Whether you plan to publish a comic or a manuscript, you need to manage a volume of material. Enter Celtx, which combines a database and text editor into a comprehensive tool. **BY ANDREAS KNEIB**

When you are ready to sit down and write that book you always think about, you might want some help getting organized. Here, I will examine Celtx [1], a tool that helps you compose comics, radio dramas, movie scripts, novels, or plays. (See also the “Alternatives to Celtx” box.)

Installation

Installing Celtx is easy, and it works right out of the box. In the download section of the Celtx project page [2], program packages are available for Windows, Mac, and Linux.

After downloading Celtx, you can unpack the tarball in the shell by typing `tar xvzf Celtx.tar.gz`, then type `cd celtx` to change to the newly create directory. When you get there, type `./celtx` to launch the writing tool.

Getting Started

On launch, Celtx comes up with a clear-cut splash screen (Figure 2) that gives you a couple of options. On the left side of the window, Celtx offers a selection of project templates that cover topics such

as text, movie, theater, comic, and radio play. The right side of the splash screen offers a selection of classical works prepared as examples – such as the H. G. Wells’ radio play *War of the Worlds* [5] or Oscar Wilde’s *The Importance of Being Earnest* [6] – that will give you an excellent overview of what Celtx offers.

After selecting Wilde’s play, you will see a number of tabs in the left-hand sidebar. The tabs give you a clear overview of the wide range of concerns involved in producing a play. For example, the *Stageplay* tab contains the literary work with acts, setting descriptions, and dialogs in a text editor.

The editor might seem spartan, but it works in an intelligent way with the use of conventions and indenting that distinguish a play or other literary scripts in printed format. This feature distinguishes Celtx from Emacs or other word processing software. For example, the software attempts to identify the format of any text you type as a *Dialog*, *Heading*, or *Figure*. If the editor fails to identify the format, you can choose the correct format from a pull-down menu on the left-hand side. Also, the editor has a function for converting the complete script, including the cover, to PDF format (see Figure 3).

Alternatives to Celtx

Besides Celtx, a number of other programs take a similar approach to facilitating the writer’s daily grind. If you prefer a more spartan approach, StorYBook [3] – a Java program – is a good choice, providing a good database that offers views and options to suit your individual requirements.

Amateur writers with serious ambitions might want to check out Writer’s Café [4]. Installation is simple, and the software integrates seamlessly with the Gnome desktop. Writer’s Café works with OpenOffice, is just as feature rich as Celtx, and offers many helpers and rules for users who are interested in learning the writer’s trade from scratch. In contrast to its competitor, only the demo version is free; the full version of Writer’s Café will set you back about US\$ 65 (EUR 41, UK£ 33).

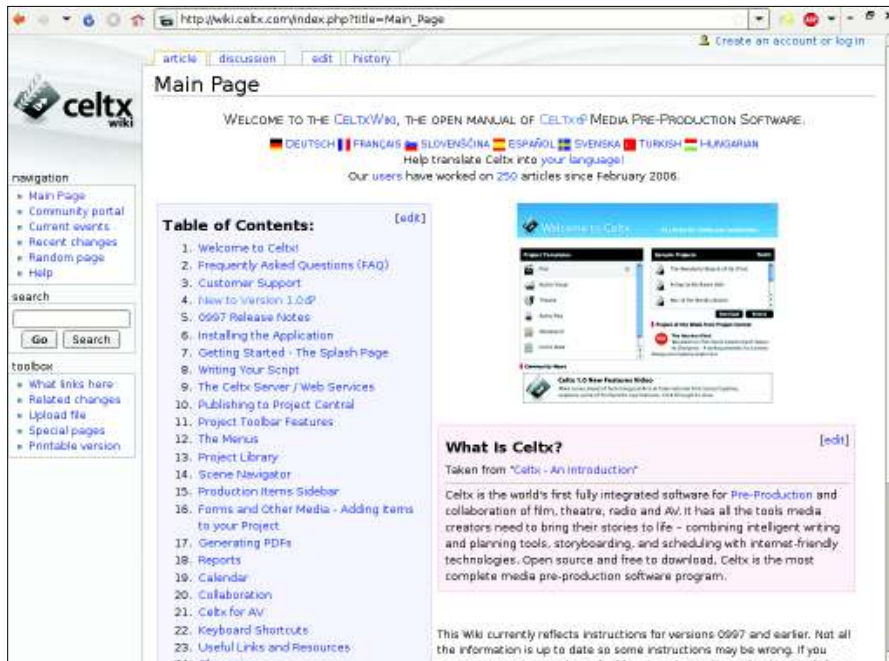


Figure 1: If you need help, check out the wiki on the Celtx website.

At first, the *Tea Images* tab, with its pictures of teacups and sandwiches, might seem slightly strange. In fact, this is a suggestion for the set in scene 1 (Figure 4). With this option, you can add sketches or photos to a project, scene by scene, and access this data whenever needed. The *Production Schedule* tab takes you to a calendar that lets you plan the production down to the minute. Finally, the *Master Catalog* gives you a database with resources such as photos and addresses of actors involved, locations, and set designs – all of which are organized in a clear-cut way, scene by scene. The next section tells you how to work with these resources.

The sidebar on the left of the window offers a further option for breaking down your material. In the case of Wilde's

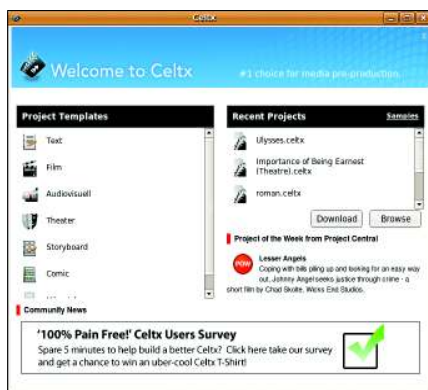


Figure 2: Celtx displays a splash screen that offers a selection of project templates and examples.

play, you will see the *Project Library* in the top half. This contains a filesystem tree with the root of the project named *Importance of Being Earnest* containing various entries. The entries include tabs such as *Master Catalog*, *Stageplay*, *Character Catalog*, *Synopsis and Summary*, and *Oscar Bio*, a short biography of Oscar Wilde (see Figure 5).

The second part of the sidebar changes its appearance or is hidden – depending on the context of the project element you are accessing. For example,

if you click the production equipment in *Tea Images*, you will see the play's *Sequences* in the lower half of the sidebar. The *Stageplay* view breaks the play down into acts and scenes. The lower part of the sidebar helps you navigate quickly through the acts and scenes via the view shown in the right pane.

Of course, you do not need to keep the tab names and descriptions – you can name the project library, the files, and the tabs to suit your own needs.

Authoring Time

Not every piece of work will be as large as this sample. Although Celtx is used mainly for drafting theater plays; scripts for radio, movies, and TV; and comics, the program can also help you structure a longer piece of prose, such as a novella or novel. As a simple example, imagine you are planning a novel titled “The Magic Mountain” [7]. First you would choose the *Text* template on the Celtx welcome window. Celtx would then come up with the *Text* and *Main catalog* tabs, which you would see in the sidebar below *Project Library*.

The next step would be to assign a name to your project. To do so, just right-click the root directory of the project tree in the side bar, which is originally titled *Unnamed*. In the context menu, select *Rename* and enter “The Magic Mountain” in the dialog that then appears.

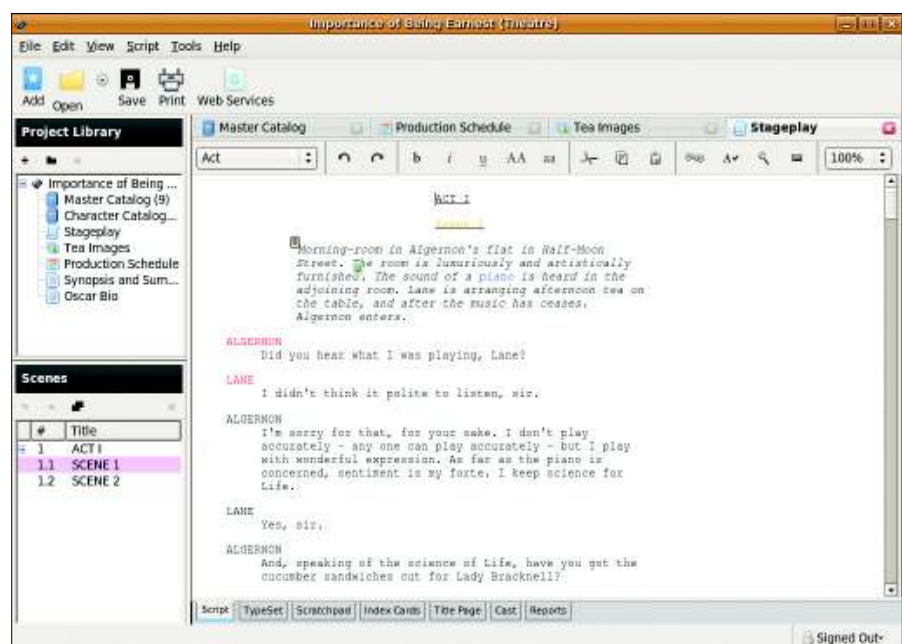


Figure 3: The program's text editor offers a number of extras, such as automatic format identification and a PDF converter.



Figure 4: If you are planning a play, you can collect the set design elements in the Tea Images folder.




Figure 5: The sample Oscar Wilde play in Celtx includes a biography of the Irish author.

Clicking the *Main Catalog* tab opens a list with the headings *Name*, *Category*, *Description*, *Tags*, and *Media*. The + *Add* button lets you add elements to the empty list, such as *Characters*, *Scene Details*, or *Location*. The element list contains all kinds of entries, but you might be able to do without fields like *Costumes* or *Stunts* for a novel.

Because a good story depends greatly on powerful characters, you will want to select *Characters* and think of a name, which you then type into the text box in the element window – say *Hans Castorp*. The right window now tiles horizontally with the tab *Character: Hans Castorp*. In the fields, you can enter this character's characteristics.

The field labels are self-explanatory. Here you can add an image to the *Media* field – a portrait of the character perhaps. The *Detailed description of body*, *Motivation*, and *Background to character* fields round off the image of your novel character.




Figure 6: Compiling the story scene by scene.

To complete the *Characteristics*, you also need to consider the character's main role in your story: this could be the *Protagonist*, *Antagonist*, *Best Friend*, or something completely different.

At this point, you can add one character after another to your narrative, as well as multiple *Locations*; however, if you click *Map of scene* in the *Address* section, the program launches your web browser and takes you to Google Maps.

Adding a *Scene details* element lets you define precisely what happens in which scene, who the protagonists and antagonists are, what their motivations are, and whether they reach these objectives in the scene. This element is important to a story, in that good material always has scope for conflict.

To create a storyboard, including a short description of all its sequences, click the plus sign (+) in the sidebar or click the *Add* icon in the taskbar. Again, you can choose from a list of elements that you might recognize from your first steps with the Oscar Wilde example. In this case, you need the *Storyboard*. Press the *Sequence* button to add one sequence after another (Figure 6). For example, 1. *Hans Castorp takes the train from Hamburg to Davos*, 2. *Hans Castorp arrives in the sanatorium in Davos*, and so on, until your story unfolds in front of your eyes – all you need to do now is write it.

But That's Not All, Folks ...

To take a break from your work, select *Save Project* in the *File* menu at the top of the Celtx window. If you are collaborating on the project with various artists

or authors, you can upload your co-production to a Celtx server. The program offers a wizard to guide you through the formalities of logging on. To reach the service, click the *Web Services* button in the tool's taskbar. In addition, Celtx offers a wiki, a forum, and a video tutorial. To access these resources, launch the matching options with your web browser in the *Help* menu.

Conclusions

Celtx is a feature-rich program, but despite this, the learning curve is not too steep: Computer newcomers should have no trouble navigating around the program, thanks to the intuitive interface. Of course, Celtx can't take one responsibility off your shoulders – that of being creative and brave enough to put your own ideas down on paper. ■

INFO

- [1] Celtx project homepage: <http://celtx.com/>
- [2] Celtx download page: <http://celtx.com/download.html>
- [3] StorYBook website: <http://storybook.intertec.ch/>
- [4] Writer's Café: <http://www.writerscafe.co.uk/>
- [5] War of the Worlds: [http://en.wikipedia.org/wiki/The_War_of_the_Worlds_\(radio\)](http://en.wikipedia.org/wiki/The_War_of_the_Worlds_(radio))
- [6] The Importance of Being Earnest: http://en.wikipedia.org/wiki/The_Importance_of_Being_Earnest
- [7] Mann, Thomas. *The Magic Mountain*. 1929. http://en.wikipedia.org/wiki/The_Magic_Mountain

Controlling multiple systems simultaneously with Synergy

LATERAL CONTROL



Synergy lets you combine multiple PCs into a single virtual desktop.

BY FLORIAN EFFENBERGER

Synergy is a small but powerful tool that lets the user control multiple PCs with a single keyboard and mouse – even across operating system boundaries. To run Synergy, you'll need at least two PCs, each with its own operating system, monitor, and working network adapter. The software supports Windows from 95 through Vista, Mac OS X version 10.2 or newer, and not only Ubuntu but other Linux distributions with the latest X server. Prebuilt packages for Windows and Mac OS X are available from the Synergy homepage [1]. Ubuntu offers a Synergy package in its Universe repository. For RPM-based Linux systems, an RPM package is also available.

First imagine a typical network administrator's workplace with a professional

desktop running on Ubuntu and a small notebook running on Vista next to it. To remove the need to keep swapping between keyboards, the administrator decides to install Synergy. Most of the work will be done on the desktop, the Ubuntu system. In Synergy-speak this master system is known as the *server*, and the

administrator will be using the keyboard and mouse attached to it. In line with this, all other devices are *clients*.

Configuration

Before you start using Synergy, you need to configure the program by editing the `/etc/synergy.conf` or `~/.synergy.conf` text file. The program relies on the concept of screens: Every computer in the group, whether it is a server or a client, is a screen with a precisely defined position – just like assigning positions to applications on a multi-display setup. Therefore, you must add the screen name, its aliases, and the position relative to the other devices to the configuration file. It is important to define the position in both directions, left and right. The setup for the case study here, and comments

Enabling Universe

Ubuntu organizes its software packages in numerous repositories. *Universe* is a hold-all for packages that are not as comprehensively supported and maintained as some others. To use *Universe*, you first need to enable the entry in the `/etc/apt/sources.list` by removing the pound sign. Then, after running `apt-get update`, you can install various new packages, including Synergy.

on it, are given in Listing 1. Many more options are available, all of which are documented on the Synergy configuration page [2].

A word of warning is appropriate here: Make sure you enter all the options in the configuration file in lowercase. In addition, pay attention to line breaks because Synergy is finicky and will not load the file if you make a mistake. After you complete the setup, launch the Synergy server on Ubuntu as a normal user by typing `synergys`. Adding `-f` prevents Synergy from running in the background as a service.

QuickSynergy

QuickSynergy provides an even more convenient approach to configuring the tool. On Ubuntu, you can download the QuickSynergy package from the *Universe* repository and install it, then launch it via *Applications | Tools | QuickSynergy*. Unfortunately, the program failed to launch a working server when I tried it in our lab.

The Vista client – that is, the PC I will be controlling via the Ubuntu system – is easier to configure. After completing the install, you can launch Synergy directly via the Start menu, which takes you to a neat-looking graphical interface (Figure 1). To connect to the server, select the *Use another computer's shared keyboard and mouse* option and enter the name of the server machine. Professionals might also be interested in setting additional options, such as the logging level, auto-start, and network details.

Linking Screens

After configuring your clients, just type `synergyc server IP` (or click *Start* if you are working on Windows) to connect the screens. At first, this is all fairly spectacular – you can use both systems in the normal way.

But if you move the mouse cursor over the right-hand border of the screen on the Ubuntu desktop, it shifts to the Vista desktop, just as on a multi-display sys-

tem – the difference is that you are using two operating system platforms. Keyboard input will also reach the client while the focus is on it (i.e., the mouse cursor is on the client screen). In addition, Synergy synchronizes the clipboards between the two systems; according to the developers, this feature automatically sets the correct character set and converts line breaks between operating systems – a perfect way to copy longer texts and configuration files centrally. Pressing *Scroll* disables Synergy temporarily if needed.

Many more options are available for you to set the configuration file. Among other things, Synergy supports keyboard mapping between the server and the clients, lets you configure screen areas where no screen paging will take place, and lets you run certain actions with the press of a button.

Of course, nobody's perfect, and Synergy failed on a couple of counts in our lab. For example, it failed to synchronize the screensavers and centrally lock all the screens.

Listing 1: Case Study

```
01 # Define screens
02 section: screens
03  ubuntu:
04  vista:
05 end
06
07 # Alternative names
08 section: aliases
09 # ubuntu -> desktop
10  ubuntu:
11    desktop
12 # vista -> notebook
13  vista:
14    notebook
15 end
16
17 # Screen positions
18 section: left
19 # vista: right of ubuntu
20  ubuntu:
21    right = vista
22 # ubuntu: left of vista
23  vista:
24    left = ubuntu
25 end
```

Security Note

As Synergy's authors point out on their homepage, Synergy does not possess anything in the line of authentication or encryption [3]. To be on the safe side, you might want to set up an SSH tunnel and encrypt all your data.

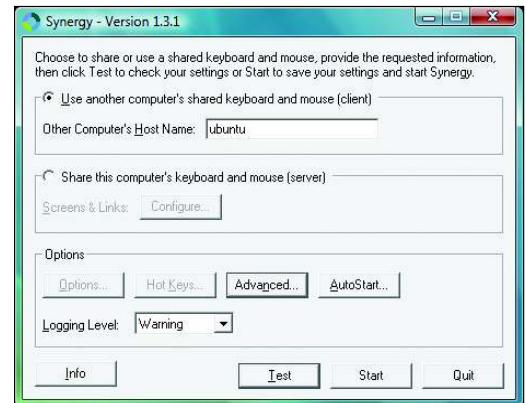


Figure 1: Synergy as a client on Windows Vista.

According to the homepage, the Mac OS X variant of the program in particular is not as mature as the Linux and Windows versions.

Conclusions

Synergy offers an interesting approach to managing multiple computers centrally without investing in additional hardware. In contrast to legacy solutions, every machine keeps its own display. The program is definitely a boon for anybody who needs to run several machines, and Synergy's cross-operating system clipboard support is a really useful feature that removes the need to copy text files. However, if I had three wishes, one of them would be easier configuration on Linux. ■

INFO

- [1] Synergy homepage: <http://synergy2.sourceforge.net>
- [2] Configuration options: <http://synergy2.sourceforge.net/configuration.html>
- [3] Security notes: <http://synergy2.sourceforge.net/security.html>

THE AUTHOR

Florian Effenberger has been a free software evangelist for many years. He is the Co-Lead of OpenOffice.org's international marketing project and a member of the board of OpenOffice.org Deutschland e.V., a German NGO. His work mainly focuses on designing enterprise and school networks and software distribution solutions based on free software. Florian is a regular contributor to various German and English language publications, in which he investigates, among other topics, legal issues.

Make yourself at home

ENVIRONMENTAL



Galyna Andushko, 123rf.com

Experienced users on Unix/Linux-based systems such as Ubuntu sometimes prefer to work in text mode from a terminal window. We'll show you some terminal tricks with environment variables. **BY HEIKE JURZIK**

Ubuntu's Terminal accessory opens a text-based command window that uses a command system called Bash (Bourne Again SHell).

Listing 1: env Example

```
01 huhn@asteroid:~$ env
02 TERM=xterm
03 SHELL=/bin/bash
04 USER=huhn
05 LS_COLORS=no=00:fi=00:...
06 PATH=/usr/local/bin:/usr/bin:/bin:/usr/bin/X11:/usr/games:/home/huhn/bin
07 IRCSERVER=irc.freenode.net
08 LANG=de_DE@euro
09 ...
```

To access the terminal window, select *Accessories | Terminal* in the Ubuntu main window.

If you enter a command like *ls* in the shell, the shell knows you want to call the */bin/ls* program. Bash knows where to look for executables, so you don't need to type in the full path. This path is defined in what is known as an environment variable. Variables of this kind are also used to modify the appearance of the shell prompt, set the time zone and

the user's home directory, and many other things.

An environment variable has three parts: the name, an assignment operator (=), and a value:

```
variable=value
```

The name can include upper- and lowercase letters, underlines, and numbers. The only restriction is that a name can't start with a number. If the value of the

Tip

Instead of using two separate commands (e.g., *LESS="-X"; export LESS*), you can define and export all at once: *export LESS="-X"*.

GLOSSARY

builtin: an abbreviation for "built-in command." Commands like this are built in to the shell, and you do not need to run a special program to use them.

GET YOUR HANDS ON SOME HOT NEW BOOKS FROM APRESS

Listing 2: export Example

```
01 huhn@asteroid:~$ export
02 declare -x CHARSET="latin1"
03 declare -x LS_COLORS="no=00:fi=00:..."
04 declare -x IRCSERVER="irc.freenode.net"
05 ...
06 huhn@asteroid:~$ export -n IRCSERVER
```

variable contains blanks or nonstandard characters, you have to double-quote your input:

```
LESS="-X"
```

Welcome Home

Every user on an Ubuntu system has a place he or she can call home; Bash provides a fully furnished environment for every single account. To discover which variables your home uses, type *env* (for “environment”).

Listing 1 shows an example. Some of the variables in Listing 1 are set by the user. Other shell variables are predefined but can be modified as necessary. The convention is to use uppercase letters.

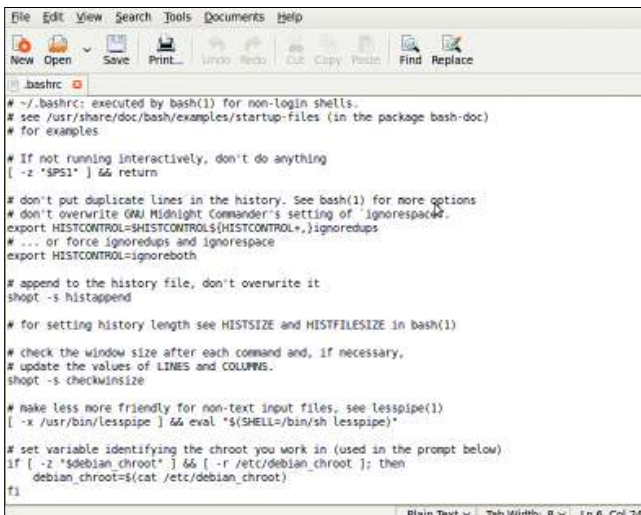
Table 1 gives an overview of common variables for Bash.

Looking into Cupboards

The *echo* command allows you to discover the value of a variable. To display the value of an environment variable, use the *echo* command with a dollar sign in front of the variable name:

```
huhn@asteroid:echo $LANG
de_DE@euro
huhn@asteroid:echo $PS1
\u@\h:\w\$\
```

The shell replaces the variables *\$LANG* and *\$PS1* with their respective values and then calls *echo* to output the text.



```
File Edit View Search Tools Documents Help
New Open Save Print... Undo Redo Cut Copy Paste Find Replace
~/.bashrc
# ~/.bashrc: executed by bash(1) for non-login shells.
# see /usr/share/doc/bash/examples/startup-files (in the package bash-doc)
# for examples

# If not running interactively, don't do anything
[ -z "$PS1" ] && return

# don't put duplicate lines in the history. See bash(1) for more options
# don't overwrite GNU Midnight Commander's setting of 'ignore_space'.
export HISTCONTROL=$HISTCONTROL,:ignoredupes:
# ... or force ignoredupes and ignorespace
export HISTCONTROL=ignoreboth


# append to the history file, don't overwrite it
shopt -s histappend

# for setting history length see HISTSIZE and HISTFILESIZE in bash(1)
# check the window size after each command and, if necessary,
# update the values of LINES and COLUMNS.
shopt -s checkwinsize

# make less more friendly for non-text input files, see lesspipe(1)
[ -x /usr/bin/lesspipe ] && eval "$(SHELL=/bin/sh lesspipe)"

# set variable identifying the chroot you work in (used in the prompt below)
if [ -z "$debian_chroot" ] && [ -r /etc/debian_chroot ]; then
  debian_chroot=$(cat /etc/debian_chroot)
fi
```


Figure 1: The *.bashrc* file in your home directory provides information on the Bash environment. First, choose Places in the Ubuntu main window and then Home Folder. If you don't see *.bashrc* in the File Browser, select the View menu and click on Show Hidden Files.




Peter Seebach
978-1-4302-1043-6
\$34.99 | 376 pp | November 2008



Ron Peters
978-1-4302-1841-8
\$39.99 | 330 pp | December 2008




Kirk Bauer and Nathan Campi
978-1-4302-1059-7
\$49.99 | 425 pp | December 2008



Andy Channelle
978-1-4302-1590-5
\$39.99 | 500 pp | December 2008



Sander van Vugt
978-1-4302-1622-3
\$44.99 | 400 pp | December 2008



Sander van Vugt
978-1-4302-1082-5
\$39.99 | 424 pp | September 2008

For more information about Apress titles,
please visit www.apress.com

Don't want to wait for the printed book?
Order the eBook now at
<http://eBookshop.apress.com!>

Apress[®]
THE EXPERT'S VOICE™

Language and Locale

The `LANG` command lists the language and locale settings for your environment. To see the locale options available on your system, type:

```
locale -a
```

If your Ubuntu system is configured for US English (the `en_US.utf8` locale) but your preference is for British English, you can change the language setting with the following command:

```
LANG=GB
```

The shell even lets you define the language setting for certain specific applications:

```
LANG=GB firefox
```

Also, you can configure your environment for different languages:

```
LANG=FR firefox
```

If the language you want to add doesn't appear with the `locale -a` command, you will need to install it through the Ubuntu repositories. To do so, try going to the Synaptic package manager

and looking for your language with the Search tool.

Change

As you have seen, you can easily change a variable by entering its name, the assignment operator, and the value. Variables defined in this way are only valid for the current terminal session. To export the settings, you need the aptly named `export` command. This command makes the settings accessible to child processes and subshells.

For example, if you are an IRC user, you can set the `IRC_NICK`, `IRC_NAME`, and `IRC_SERVER` variables for the current Bash shell to tell command line--based IRC clients which nick and name to use to log on automatically to a specific IRC server. If you then launch a second shell from the first one and try to output the values of the new variables, you won't see a thing. First you need to quit the second shell by typing `exit` (Ctrl + D) and export the new variables. Now open a new shell from the first one and you will see that it has inherited variables, which are now available in the subshell.

The `export` builtin has more tricks up its sleeve. In combination with the `-p` option, or when called without any parameters, it displays a list of all exported

variables. The `-n` option deletes a variable from the list. Listing 2 gives a couple of examples.

Home Improvements

Any variables you *export* are only effective in the current shell and its child processes. To set up a variable permanently, you need to modify your Bash configuration file and reload your environment. To do so, add your new `export` commands to one of the Bash startup files, such as `~/bashrc` (Figure 1), then parse the file once you've finished your home improvements:

```
source ~/.bashrc
```

Listing 3 shows a couple of practical environment variables, including comments, that should make you feel at home. ■

Listing 3: Environment Variables

```
01 # prevents "less" from leaving an
02 # screen at the end of the output:
03
04 export LESS="-X"
05
06 # adds a touch of color: user
07 # prompt in
08
09 if [ $(id -u) = 0 ] ; then
10     COLOR1='\[\033[00;31m\'
11 else
12     COLOR1='\[\033[01;32m\'
13 fi
14
15 COLOR2='\[\033[00;33m\'
16 COLOR3='\[\033[00;37m\'
17 PS1=$COLOR1'\[\u@\h'$COLOR2'
18 \W'$COLOR1']'$COLOR3'\$ '
19
20 # I want the CET time zone on my
21 # server in Canada:
22 export TZ=CET
23
24 # I want ispell to use the correct
25 # character set
26 # and the correct dictionary:
27 export DICTIONARY=ngerman
28 export CHARSET=latin1
```

Table 1: Standard Bash Environment Variables

Variable	Meaning
<code>CDPATH</code>	Search path for the <code>cd</code> command
<code>EDITOR</code>	Standard text editor
<code>HISTFILE</code>	History file (e.g., <code>~/bash_history</code>)
<code>HISTSIZE</code>	Maximum number of commands in history
<code>HOME</code>	Home directory for the user account
<code>HOSTNAME</code>	Hostname
<code>LANG</code>	Language for program output, such as the date format, etc., assuming none of the <code>LC_</code> variables (see below) have been set. With the <code>locale</code> command, you can output a listing of all defined language variables.
<code>LC_ALL</code> <code>LC_</code> variables.	Country setting, such as <code>C</code> or <code>de</code> ; this overrides <code>LANG</code> and any other <code>LC_</code> variables.
<code>LC_MESSAGES</code>	Language for program and error messages
<code>LC_TIME</code>	Time format
<code>LOGNAME</code>	Login name of user
<code>MAIL</code>	Path to the user's mailbox (incoming mail)
<code>MANPATH</code>	Search path for man pages
<code>PATH</code>	Search path for executables
<code>PS1</code>	Default appearance of the shell prompt. For non-privileged users, this is typically <code>\u@\h:\w\\$(huhn@asteroid:~\$</code> in this example); for the administrator, this is typically <code>\h:\w\\$(asteroid:~#)</code> .
<code>PWD</code>	Name of the current directory ("print working directory")
<code>SHELL</code>	Full pathname of the current shell (<code>/bin/bash</code>)
<code>TERM</code>	Terminal settings, such as <code>xterm</code> or <code>vt100</code>
<code>TZ</code>	Time zone, such as <code>CET</code> or <code>MET</code>

Discovery Guide



Peek inside for a practical introduction to Ubuntu, including tips on installing, configuring, and exploring your new Ubuntu system.

Installing Ubuntu.....	62	Printing	74	Installing Software	82
Exploring.....	64	Multimedia	76	Email	84
Networking	67	Burning CDs.....	78	Taking Notes	88
Security.....	71	Managing Photos.....	80	OpenOffice	90



Installing Ubuntu

If your hardware supports Ubuntu, you should have your new system up and running in less than an hour.

BY CARSTEN SCHNOBER

The Ubuntu DVD enclosed with this issue installs on most contemporary Intel-compatible PC systems. The machine you use will need at least 384MB of RAM. Certain graphics cards, WLAN chipsets, and other components might not be supported by Ubuntu. Included on the DVD is a Live boot option that allows you to boot Ubuntu from the DVD without disturbing the contents of your hard disk. To start, you might want to perform a Live boot from the DVD to test for any hardware compatibility issues before you do a conventional installation. If you boot to the Live system, you will find a folder called *Examples* on the desktop with a variety of test files for exploring multimedia and office applications and other Ubuntu features.

If you experience problems with hardware support, see Ubuntu's online help sources for additional information. The Ubuntu Jaunty Jackalope release notes are a good starting point [1]. Ubuntu also provides forums [2], a wiki [3], and other sources for troubleshooting information [4].

Getting Started

First, place the DVD included with this issue in your DVD drive and reboot. If all goes well, the boot menu shown in Figure 1 will appear on your screen. If you don't see the boot menu, make sure your BIOS is configured to start the system from the DVD drive. (See the box titled "Boot Order.")

Boot Order

The BIOS (Basic Input/Output System) decides which drive to use to load the operating system.

If your system doesn't boot to the DVD by default, watch for the message displayed when you power on, telling you now to change the boot order. Look for something like *Press X for boot manager*, where *X* represents a key (often the F8 key).

If your BIOS doesn't have a boot manager, you need to change the order manually. To

The first item in the boot menu, *Try Ubuntu without any changes to your computer*, lets you boot the system in Live mode from the DVD. As mentioned previously, you can try a Live boot first before committing to a hard disk install. To navigate the menu, use the arrow keys and press Enter to select. The *Check disc for defects* and *Test RAM* options take you to a couple of test routines. If you have decided that you just want to boot your old system normally and ignore the contents of the DVD, select *Boot from first hard disk*. If you are sure that you want to install Ubuntu on your hard disk, select *Install Ubuntu*.

Step by Step

Most of the prompts from the Ubuntu installer are just informational and are not necessarily

do so, enter the BIOS menu by pressing the necessary key (e.g., F2 or Del) when you power the system on. (The correct key will depend on your flavor of BIOS. See the vendor documentation for your motherboard. Instructions for reaching the boot menu often appear in the screen messages when you power your system on.)

In the BIOS menu, look for an item labeled *Boot*. Set the options so that the DVD drive has a higher priority than the hard disk.

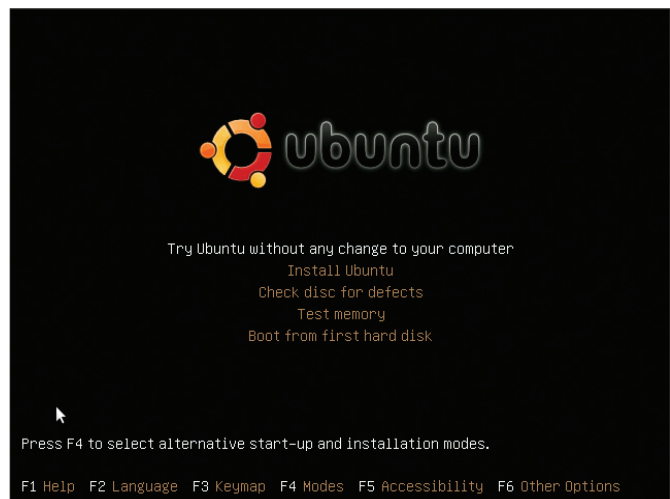



Figure 1: After booting from the installation DVD, you have a large selection of options. Take time to get to know the system before you load it onto your hard disk.

technical in nature. The first thing you need to do is specify your preferred desktop language and locale.

In the next dialog, you need to tell Ubuntu the time zone in which your computer resides. To set up your clock correctly, choose a city in the list that shares the same time zone as your machine.

The keyboard layout dialog follows. The system will suggest a keyboard layout that reflects your choice of language, and it will assume a standard PC keyboard. If needed, variants such as Macintosh keyboards or keyboards with no *Dead Keys* are available; the latter does not understand accents or non-standard characters. In the input box, you can test your choice of keyboard.

The next step lets you choose the hard disk partition on which to install Ubuntu (Figure 2).

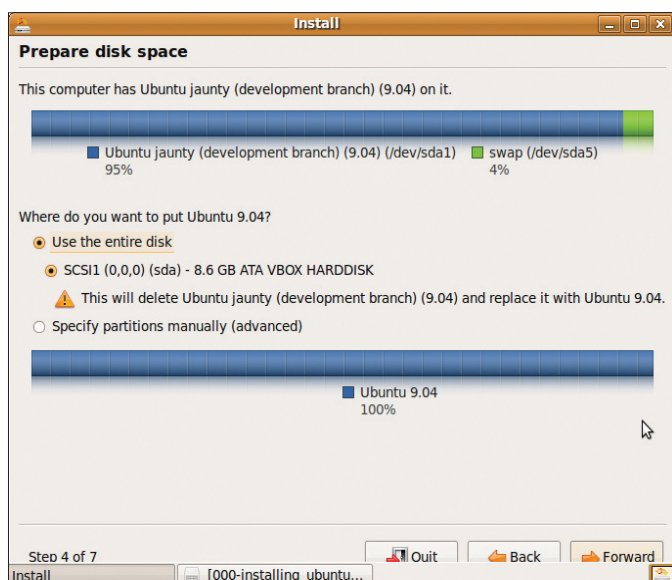


Figure 2: In the partitioning phase, you specify the hard disk area in which you will be installing Ubuntu. You also get to decide how much of the disk to use for Ubuntu.

A partition is a section of the hard disk that is treated as a single entity by the filesystem. If the disk is empty, or if you don't care about losing the data and operating systems currently residing on your disk, choose the *Use entire disk* option to use the whole disk for your Ubuntu installation.

The installer is capable of setting up Ubuntu to dual-boot with an existing Windows or Linux installation. In this case, you will need to install Ubuntu on a separate partition. All data on the partition you use to install Ubuntu will be lost. Select *Specify partitions manually* to resize, delete, and create partitions. As you can see in Figure 2, this manual configuration option is considered an advanced skill. For additional background on how to configure partitions, see the Ubuntu website [5]. Also, don't forget to back up any important files on the previous system.

Migrants Welcome!

After you decide on a partition configuration, the installer offers to parse settings and files

from existing Windows and Linux systems on the hard disk. If you have multiple user accounts, select the accounts you want to import, then select the information you want to import. The wizard will import bookmarks from Firefox and Internet Explorer, along with the wallpaper and files from the *My Files* directory structure on a Windows system. Note that this will only work if you retain the Windows system on your disk!

Finally, the installer asks you for your name, username, and password. To make life more difficult for potential attackers, the password should contain at least eight characters, some of which should be numbers and letters. The system also suggests the hostname, which is used to talk to the machine on the local network. To finish this part of your installation, you can enable the *Automatic login* option if you want to log in without entering a password when you boot the machine.

Before the installation starts, the installer shows you an overview of your configuration choices (Figure 3). Just make sure to read the

details carefully; once you confirm, the installation will start.

Let the Games Begin

After the installation, you should remove the DVD from the drive, shut down the system, and press Enter to boot your new Ubuntu system for the first time. The computer will load Ubuntu and present a login window in which you can type your username and password (unless you enabled the *Automatic Login* option earlier).

After you log in, it is a good idea to upgrade your system with any recent bug fixes or security updates that have appeared since this DVD was published. New packages are installed over the Internet, so you'll need to have a working Internet connection to download and install the latest packages.

To update your system, select *System | Administration | Update Manager* to launch the Update Manager. For more on installing software packages in Ubuntu, see the article on managing packages. ■

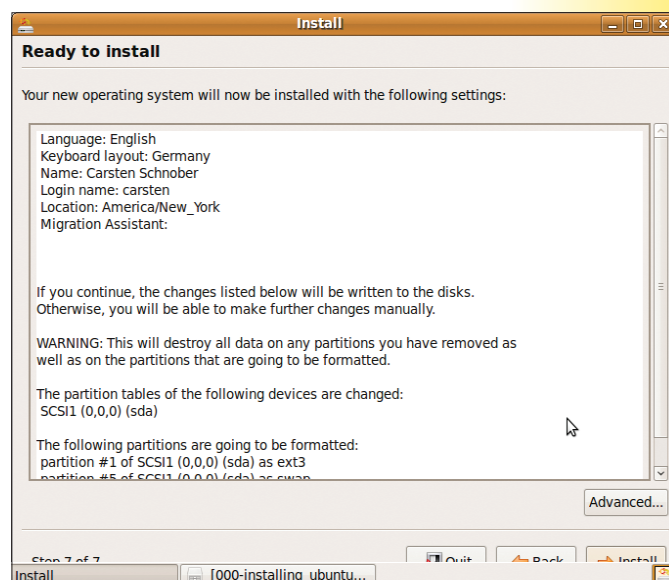


Figure 3: Carefully check the overview before you install your new operating system. Any pre-existing data on the partition to which you are installing will be lost.

Boot Options

If your system fails to boot or does not work properly with the defaults, you can set manual boot parameters to resolve the issue. In the boot menu, press F1 for an overview. If needed, you can press F2 and F3 for an alternative language and keyboard options. F4 lets you use a safe graphics mode – a useful workaround for problems caused by missing graphics adapter drivers.

Pressing F5 sets up the display for special needs. F6 lets you select options for buggy motherboards. *acpi=off* disables the use of the *Advanced Configuration and Power*

Management Interface (ACPI). Your computer's power-saving features will not work with this setting, but you might be able to convince imperfect hardware to cooperate with Ubuntu.

The *noapic* and *nolapic* options also are commonly used for imperfect motherboards. These options pass the responsibility for hardware addressing to the operating system rather than the motherboard. No simple rule lets you determine which of these options will help you solve boot issues, but I suggest trying *acpi=off*, *noapic*, and *nolapic*, in that order.

Info

- [1] Ubuntu Jaunty Jackalope release notes:
<http://www.ubuntu.com/getubuntu/releasenotes/904>
- [2] Ubuntu forums:
<http://ubuntuforums.org>
- [3] Ubuntu wiki:
<https://wiki.ubuntu.com>
- [4] Official Ubuntu documentation:
<https://help.ubuntu.com/>
- [5] How to partition:
<https://help.ubuntu.com/community/HowtoPartition>

Exploring the Desktop

Make yourself at home on the Ubuntu desktop.

BY CARSTEN SCHNOBER

Unlike Windows and Mac OS X, Ubuntu and other Linux systems separate the desktop graphical user interface (GUI) from the underlying system. The standard desktop included with the main Ubuntu release is Gnome [1]. This introduction assumes you are using the standard, Gnome-based desktop included on the DVD accompanying this magazine. Other desktops, such as KDE [2] and Xfce [3], are also popular. Kubuntu [4] is an attractive KDE-based variant of the Ubuntu distribution. If your computer does not have much RAM, you might prefer Xubuntu [5], which installs with the Xfce desktop. (See the box titled “Desktops.”)

Logging In

The Ubuntu desktop launches with the login screen (Figure 3) when you enter your username and matching password. Clicking on *Options* lets you select a different language or desktop environment.

If you don't share your computer, you can avoid the need to log in by enabling *Automatic Login*. Either enable this option during the installation or after logging in to the desktop. To

enable automatic login, click *System | Administration* and then select *Login Window* to display the Login Window Preferences dialog box. The *Security* tab takes you to the *Enable Automatic Login* function (Figure 2); you also need to select your own user account from the list.

Menus

Once you have finished logging in, you will find three menus in the top left corner of your screen (Figure 1). These menus – Applications, Places, and System – let you launch programs, navigate the filesystem and the network, and configure your system and desktop, respectively.

The Applications menu is arranged topically by program type. For example, *Universal Access* contains software to support users with special needs; the *Internet* option lets you choose a web browser, as well as Internet telephony, email, and chat programs. The *Add/Re-*

move entry lets you search for and install additional programs or remove existing programs, by checking or unchecking the boxes next to the entries (Figure 4).

The top entries in the Places menu take you to directories in your home directory – that is, to the folders containing files for your account. The *Desktop* option opens the directory for the desktop, where your desktop files are stored. All of your personal files and directories are located below your *Home Folder*. If you select *Bookmarks | Add Bookmark* in the file browser, you can use the Places menu to access the bookmarked folders directly.

If you need to investigate files outside of your home directory, you can use the *Computer* option to access the hard disk and other external media, such as CDs and USB sticks.

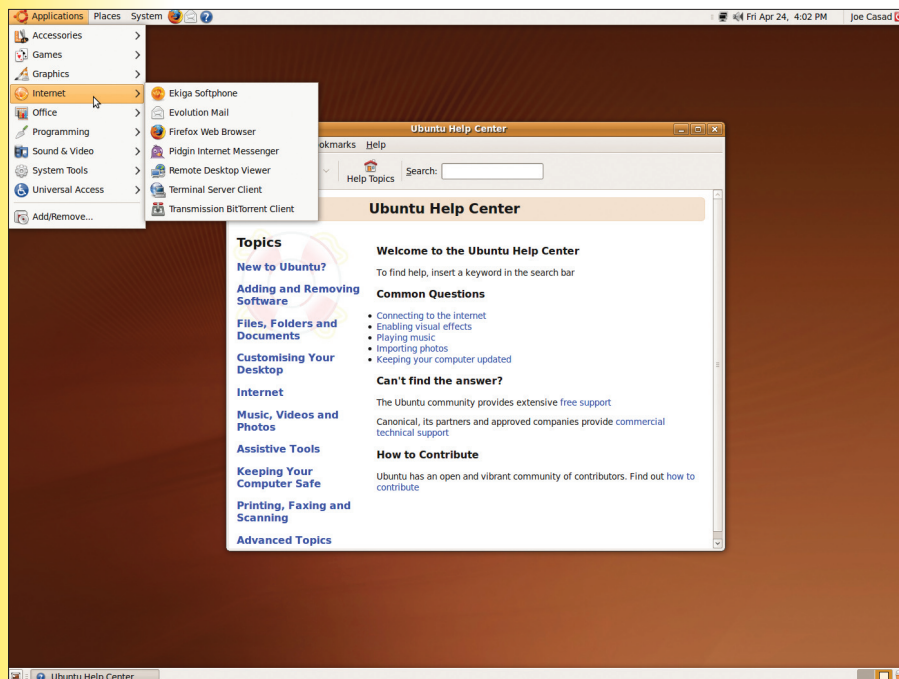


Figure 1: The menus in the upper panel give you access to Ubuntu's desktop applications.

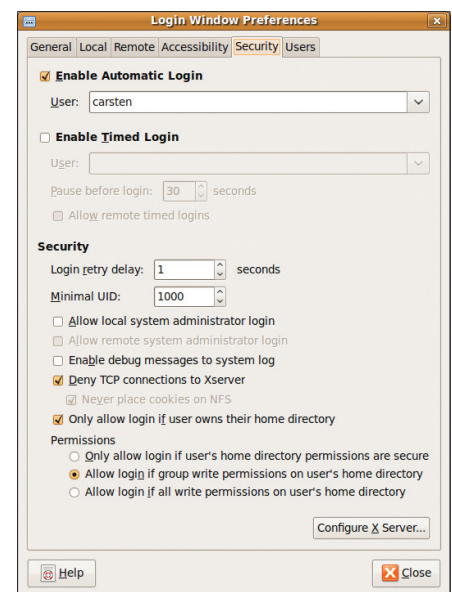


Figure 2: If you want to avoid entering your username and password, choose *Enable Automatic Login*.




Figure 3: Before you can start working, you need to log in.

The desktop opens USB sticks automatically in a separate window when you plug them in.

The *Network* entry takes you onto the local network and shows you the other machines inside your home or office network, along with any network shares. Also, you can access Internet servers directly via the desktop, although this does involve a small work-around. If you select *Connect to Server*, a window appears with support for popular transport protocols (Figure 5). To access an SSH server via the Internet, select the *SSH* entry as the *Service type*, then type the *Server* name and your *User Name*. If you want to access a particular folder after logging in, you can enter a *Folder* name. Now you're ready to click *Connect* to open the connection to the server and move your files back and forth, just as if they were on the local filesystem. To use the connection on a regular basis, you should check the *Add bookmark* option in the *Connect to Server* dialog box and type an intuitive name for the connection. Then you can use the bookmark later to open a connection.

Setting Up the System

The *System* menu lets you customize system and desktop settings. The menu items are organized in two categories *Preferences* and *Administration*. *Preferences* gives you options that do not affect the system globally. On a multiuser system, every user can choose their own settings. For example, *System | Preferences | Appearance | Visual Effects* lets you enable the 3D desktop (assuming you have 3D accelerator support). *Appearance | Display* lets you change the screen resolution. If you use an nVidia graphics card, Ubuntu will automatically call the nVidia tool to handle this task.

The other item in the *System* menu, *Administration*, is for settings that have a more fundamental effect on the system as a whole. These settings are typically available only to the system administrator. Before accessing the options in *Administration*, you will need to provide a password. (The first user account you set up in Ubuntu has administrative privileges.)

Desktops

The Ubuntu variants Kubuntu and Xubuntu install alternative desktops by default; however, you don't need a whole different Ubuntu distribution just to experiment with other desktop options. If you want to try KDE or Xfce, just launch the Synaptic package manager, look for *kubuntu-desktop* or *xubuntu-desktop*, and install one or both. However, you will need some free disk space (about 500MB). After installing the alternative desktops, boot your PC to the Ubuntu login desktop and click *Options | Select Session* (Figure 2).

MISSING LINUX MAGAZINE?



Ever have problems finding Linux Magazine on the newsstand? Just ask your local newsagent to reserve a copy of Linux Magazine for you!

Simply download our Just Ask! order form at www.linux-magazine.com/JustAsk, complete it, and take it to your local newsagent, who will reserve your copy of Linux Magazine.

Some newsagents even offer home delivery, making it even easier to ensure you don't miss an issue of Linux Magazine.



**SPECIAL SERVICE
FOR OUR UK READERS!**

www.linux-magazine.com/JustAsk

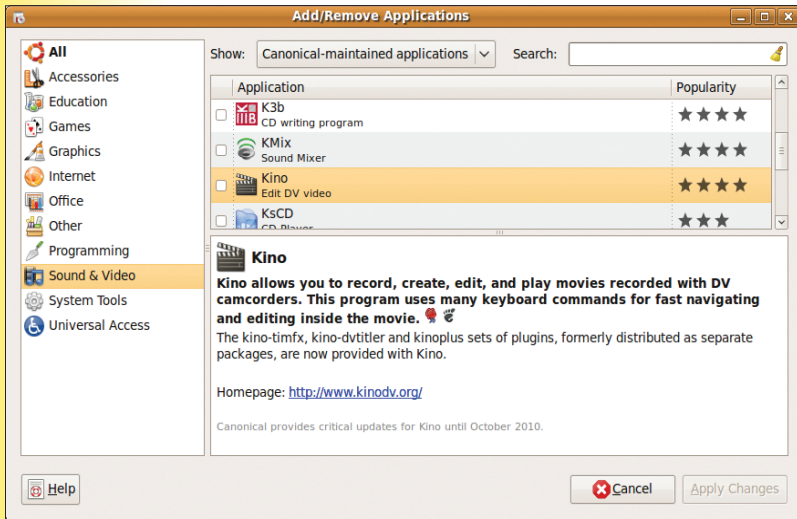


Figure 4: The Gnome menu lets you install and remove new programs at the click of a mouse.



Figure 7: Customizing the Ubuntu panel is done from the Add to Panel dialog box.

The entries in this menu item let you update your software collection by calling the Synaptic package manager, install proprietary drivers (*Hardware Drivers*), and check your hardware (*System Monitor* or *Network Tools*). The *USB Startup Disk Creator* allows you to install Ubuntu on a USB stick. This option lets you carry a bootable Ubuntu system, including your files, wherever you go.

Although the system regularly checks for updates on its own, the *Update Manager* allows you to check manually for system updates. If it finds any in automatic mode, it displays a message in the panel at the top of the screen. Clicking the icon takes you to a window in which you can automatically install the latest versions of your choice of software.

When you create a user account in *System Administration | Users and Groups*, you can set a *Profile* for each user – for example, *Administrator*, *Desktop user*, or *Unprivileged* (Figure 6). Administrators are allowed to change system settings, whereas *Desktop users* are not. Unprivileged users are given a fairly restricted workspace that even prevents access to external media and multimedia devices. The *User*

Privileges tab lets you modify the details of a user's account.

Panels

Next to the menus, you will see launchers for the Firefox web browser, the Evolution email and calendaring application, and the help system. By right-clicking and selecting *Remove from panel*, you can remove these icons from the panel. If you prefer a different email program, you can add the icon for your own mail client, instead of Evolution, by right-clicking a free space in the panel, selecting *Add to Panel*, and in the dialog that then appears, selecting *Application Launcher* (Figure 7). Now you can locate the program in the menu structure.

The desktop provides several panel applets that let you monitor system activity, integrate new functions, or just have fun. For example, you can add a *Fish* that will swim around on your desktop, make notes of events on yellow *Sticky notes*, or check the *Weather report*.

Applets are preconfigured panel elements that appear top right on your screen; they will differ slightly depending on your system. Besides the clock, you will see the network manager, and, if you have a laptop, the battery

charge status. Clicking your name lets you log off, shutdown, or restart the computer. Also, you can switch to another user account without quitting the current session. A *Guest session* is useful for one-off users because it deletes any files the guests create in their home directory when they log off again.

At the bottom edge of the screen is another panel that you can customize to suit your own needs. By default, it shows you all the program windows that are currently open. Clicking on an icon for a window moves the window to the foreground. The button on the bottom left iconizes all the active windows; the button in the bottom right is for the trash can, where any files you delete are temporarily stored until you empty the trash.

Next to the trash can in the bottom panel is the pager, which is a feature that lets you open any number of independent screens and switch back and forth between them by clicking the boxes. This setup is a useful way of separating tasks, in that you can access a fresh workspace without closing the windows on the current desktop. If the default of two workspaces is not sufficient, right-click the pager, select *Preferences*, and enter the number of *Columns* and *Rows* under the *Workspaces* field.

Newbie-Friendly

Ubuntu and the Gnome desktop offer a flexible, versatile working environment for beginners and advanced users. Other articles in this Discovery Guide provide a closer look at some tools of the Ubuntu environment. ■

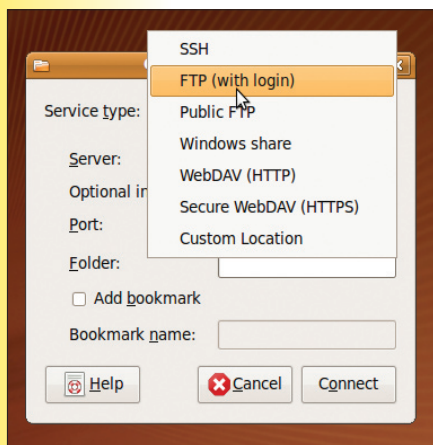


Figure 5: Choose *Connect to Server* in the Places menu to access an Internet server.



Figure 6: Setting privileges for a new user.

Info


[1] Gnome: <http://www.gnome.org>

[2] KDE: <http://www.kde.org>

[3] Xfce: <http://www.xfce.org>

[4] Kubuntu: <http://www.kubuntu.org>

[5] Xubuntu: <http://www.xubuntu.org>



Networking

Ubuntu will connect automatically if your network uses DHCP. Configuring a static connection takes a little more time. If you're going wireless, you might need to install a driver.

BY KRISTIAN KISSLING AND MARCEL HILZINGER

Ubuntu sets up networking automatically on most systems, but, depending on your configuration, you might occasionally prefer to define your settings manually. In this article, we take a look at TCP/IP networking on conventional (wired) networks as well as wireless LANs.

DHCP

On most modern PC networks, the IP address and other network configuration settings are assigned through a Dynamic Host Configuration Protocol (DHCP) server. If you are working on a home network, your home router/firewall device is probably acting as a DHCP server. If you connect directly to your ISP's network, your computer will receive an address through

your ISP's DHCP server. Your Ubuntu system is set up to receive a configuration automatically through DHCP. As long as your system is connected to a network and has access to a DHCP server, Ubuntu will configure TCP/IP networking for you.

If your system does not successfully receive an automatic connection through DHCP, make sure your computer is properly connected to the network and the DHCP server is online. Some DHCP servers assign IP addresses only to computers on a pre-authorized list of Media

Access Control (MAC) addresses. In that case, you'll need to add your computers' MAC address to the DHCP server. To find out, contact your ISP or consult your DHCP server documentation.

The MAC address is a permanent address that is associated with your network adapter. To determine the MAC address for your system, open a Terminal window – *Applications | Accessories | Terminal* – and enter *ifconfig*. Now look for the setting labeled *HWaddr* (outlined in Figure 1). Also, you can right-click on the Network Manager applet in the upper-right corner of the Ubuntu panel and select *Edit Connections*.

Configuring a Static IP Address

Before you assign a static IP address, you will need to know the address you plan to use, the subnet mask, the gateway address of the router, and the address of the DNS server. To set up a static IP address, right-click the Network Manager icon on the upper-right side of the panel

IPv4 vs. IPv6

Because “only” four billion possible IPv4 addresses exist worldwide, the plan is to replace these addresses in the future with IPv6 addresses, which comprise eight blocks of 16 bits and thus support up to 2^{128} addresses. However, in real life, IPv4 addresses are still more common.

Tip

To manually request a DHCP address from the DHCP server, open a terminal and type *sudo dhclient* (Figure 2). The server will then assign you a new IP address. This option is useful if the Network Manager isn't running or isn't co-operating.

```

carsten@carsten-desktop: /home/carsten
File Edit View Terminal Help
carsten@carsten-desktop:~ $ /sbin/ifconfig
eth0    Link encap:Ethernet HWaddr 08:00:27:bc:aa:d6
        inet addr:10.0.2.15 Bcast:10.0.2.255 Mask:255.255.0
        inet6 addr: fe80::a00:27ff:febc:aad6/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
        RX packets:98788 errors:0 dropped:0 overruns:0 frame:0
        TX packets:44547 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:146098547 (146.0 MB) TX bytes:2666052 (2.6 MB)

lo      Link encap:Local Loopback
        inet addr:127.0.0.1 Mask:255.0.0.0
        inet6 addr: ::1/128 Scope:Host
        UP LOOPBACK RUNNING MTU:16436 Metric:1
        RX packets:2 errors:0 dropped:0 overruns:0 frame:0
        TX packets:2 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:168 (168.0 B) TX bytes:168 (168.0 B)

pan0    Link encap:Ethernet HWaddr 96:4f:7d:a8:7c:cd
        inet6 addr: fe80::944f:7dff:fea8:7ccd/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:29 errors:0 dropped:0 overruns:0 carrier:0

```

Figure 1: The *ifconfig* command provides a summary of network configuration information. The *HWaddr* setting refers to the MAC address.

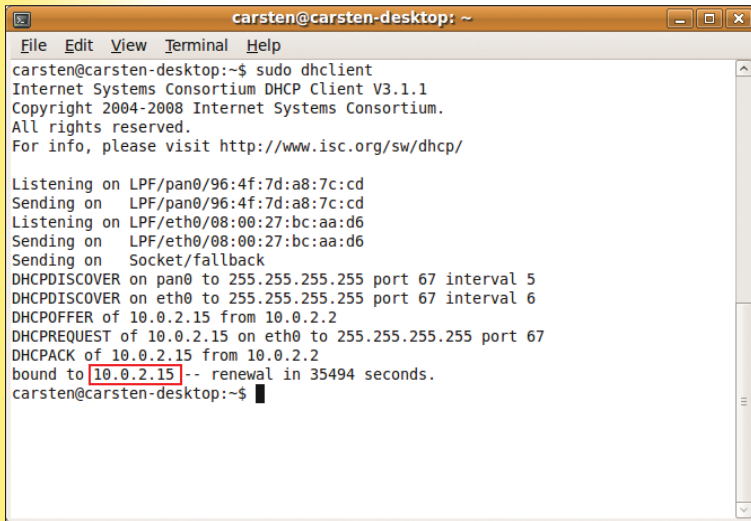


Figure 2: The `dhclient` lets you manually request a private IP address from the DHCP server.

and select *Edit Connections*. In the network connections window, click on the *Wired* tab and press the *Add* button. Start by typing a name for the connection (say, *static*), then go to the *IPv4 Settings* tab (Figure 3) and select the *Manual* entry in the drop-down menu. Now click *Add* and type a static IP address for the computer (*10.0.2.253* in Figure 3), followed by the netmask (*255.255.255.0*) and the gateway (i.e., router) address (*10.0.2.2*). Because the router on a small network typically is also the DNS server, the server address in Figure 3 is also *10.0.2.2*. Clicking *Apply* saves the changes. To enable the new settings, left-click the Network Manager icon and select the *static* connection.

Network Problems

The `ifconfig` command provides a summary of the IP configuration. Also, you can right-click the Network Manager icon and choose *Con-*

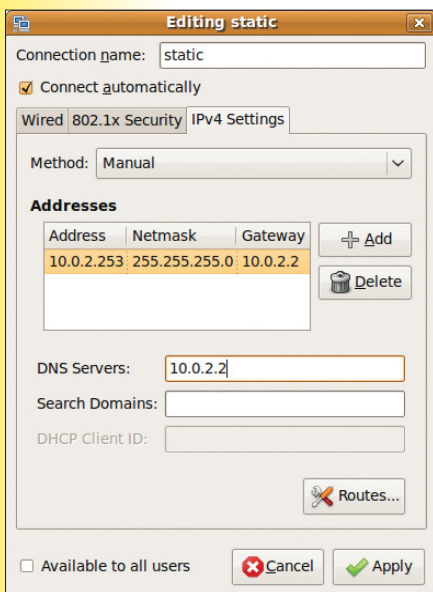


Figure 3: Setting up a network connection with a static IP address.

nection Information for a summary of network settings.

If you can't access the Internet, start by pinging the router and its gateway address:

```
ping 10.0.2.2
```

(Type `Ctrl+C` to stop the output and return to the command prompt.) If this does not work, the router might not be working or the cable between the router and your computer might be damaged. Also, you can ping the "loopback address" *127.0.0.1*, which gives an indication of whether your networking software is successfully connected to your the networking hardware.

If your connection to the router is working, you can try pinging an IP address that you are familiar with of a website beyond the router,

such as *80.237.227.149* for Linux Magazine Online.

If the website responds to the *ping*, the problem might be due to DNS name resolution rather than IP networking. Try pinging the same address by its domain name:

```
ping www.linux-magazine.com
```

If you can reach the address by its IP address but not by its domain name, the problem is probably with DNS.

Network Tools

Ubuntu provides a number of classic Unix/Linux configuration and troubleshooting utilities via a convenient GUI interface. To reach the Network Tools application, select *System | Administration | Network Tools* (Figure 4).

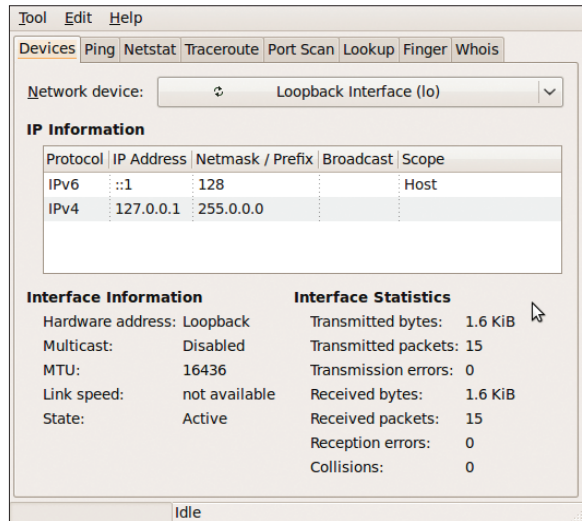


Figure 4: The Network Tools application provides a GUI interface for classic Unix network utilities.

Adding a Static Address Manually

Of course, Linux wouldn't be Linux if you couldn't delve deeper and add the IP address manually to the `/etc/network/interfaces` file. To start, stop the existing network by typing `sudo /etc/init.d/networking stop`. Then press `Alt+F2` and type `gksu gedit` to launch the Gedit editor with administrative privileges. Add the parameters from the Static IP box to the `/etc/network/interfaces` file:

```

auto eth0
iface eth0 inet static
address 10.0.2.253
netmask 255.255.255.0
gateway 10.0.2.2
    
```

The first two lines auto-enable the network interface (*auto eth0*) and tell the computer to set up a static IP address (*static*). If the computer has multiple network adapters, you need to specify the name of the adapter (*eth1*, *eth2*, etc.). The third line specifies the value for the static IP address. The fourth line contains the subnet mask (*255.255.255.0*), and the last line contains your router's gateway address, which you discovered previously. Now save the file and then open the `/etc/resolv.conf` file. This file is very important because it contains the address of the name server that converts IP addresses to domain names. Next, add the line `nameserver 10.0.2.2`, save the file, then type `sudo /etc/init.d/networking start` to restart the network. When you type `/sbin/ifconfig` in the terminal window, you should see your static IP address.

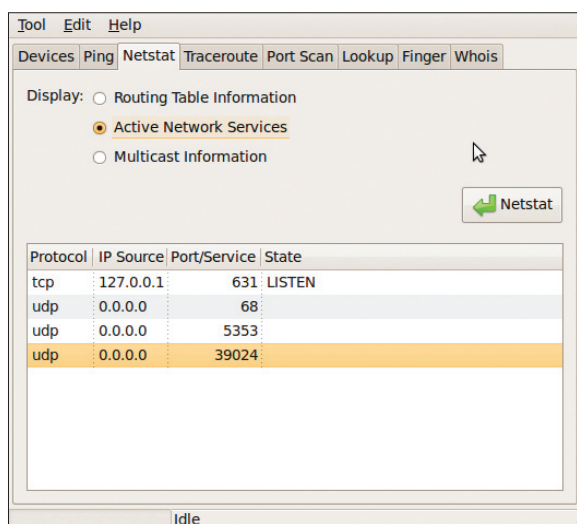


Figure 5: Studying the network with the Network Tools *Netstat* tab.

The *Devices* tab (shown in Figure 4) provides statistics on network devices. To review transmission statistics, choose a network interface in the drop-down menu. The *Ping*, *Netstat*, and *Traceroute* tabs offer a GUI approach to these classic tools. *Traceroute* shows the path through a chain of routers to a specific network address. The *Netstat* utility provides some useful network statistics (Figure 5). To update the display, choose a radio button and click the *Netstat* button.

Network Manager

The Network Manager application in the upper-right side of the panel is a central interface for viewing and editing the network configuration on your computer. Network Manager allows you to set up your network, check on the network status, and view various network statistics.

Installing the Rt73 Driver

If you use one of the very popular Rt73 chipsets on Ubuntu 9.04, the installation is very simple. Start by checking which chipset you have by typing `tail -f /var/log/messages` in a terminal to read the system log, and then plug in the stick. The output should look something like Figure 8; the `usbcore: registered new interface` line should give you the driver name used by the WLAN card. If this name is `rt73usb`, as in the figure, you can go on to install the `rt73-common` package with the firmware for the chip. Then use Network Manager to associate with the access point. If you need a different driver, search for the name of the driver in Synaptic, or check the excellent Ubuntu community help page [1].

Advanced networking technologies, such as virtual private networks, mobile broadband, and DSL, are also managed through the Network Manager. If you are looking for information on a network connection, select the connection in the Network Manager and click the *Edit* button.

Wireless Networking

Assuming the kernel supports your WLAN card, Ubuntu 9.04 will automatically launch Network Manager or a corresponding Gnome applet. In the panel, you will see a computer icon or – while the connection is opening – two gray dots that turn green if the connection attempt is successful.

If wireless support is configured correctly, clicking on the Network Manager icon reveals a list of available WLAN networks (Figure 6). When you select an entry, the applet attempts to associate with the access point for the network. If the WLAN uses encryption, a password dialog appears (Figure 7), from which you can choose the encryption type: WEP or WPA. Here, you type the password and click *Connect*. After a couple of seconds (assuming all goes well), the WLAN connection opens: The icon changes into a bar chart that shows the reception strength. If things work perfectly, the whole process should take no longer than a couple of seconds, and it will not cost you more than a couple of mouse clicks.

Wireless Drivers

Even if you don't have a cutting edge device, chances are your WLAN adapter will work



Figure 6: Left-clicking Network Manager displays the names of the WLANs within range.



Figure 7: Network Manager has identified a network with the highly insecure WEP encryption and is now displaying the corresponding interface.

with Ubuntu. The Linux kernel has built-in support for many wireless drivers, such as drivers for Intel wireless cards. However, certain devices can cause complications. If your WLAN system isn't working, one common cause is the WLAN driver.

Most USB devices use a Ralink chipset, and the current Ubuntu has the necessary Rt73 support (see the box titled "Installing the Rt73 Driver").

Linux drivers are available for most of the current crop of "Draft N" USB sticks, which use the 802.11n standard, because these sticks typically rely on the Ralink RT2870 chip, for which the manufacturer offers Linux drivers.

The Router

Home and office wireless networks typically operate through some form of wireless access point device that acts as a router. Wireless router setup normally involves accessing a web interface in a browser by typing the router's IP address. Different routers provide a range of configuration options. See the documentation for your wireless WLAN device for more on setting up your wireless router.

You should set up a name (or ESSID) for your wireless network. You might also have an option to set the frequency on which the router transmits; however, you should not change the frequency unless you are working in an environment with a large number of wireless networks that have experienced problems.

When opening a connection, Ubuntu automatically uses the channel offered to it by the access point. Depending on what country you are in, some frequencies might be restricted for private use.

Another important consideration for your wireless network is encryption. An unencrypted WLAN network can be used for Internet access by any computer in a range of about 100 meters (depending on the router's transmitter power), which means that an intruder in a nearby house or car can easily snoop on your unencrypted communication.

Many routers offer the option of WEP or WPA encryption. WEP is generally regarded as totally insecure, and thus it typically isn't used. WPA offers much better protection.

```

kkissling@kkissling-laptop: ~
Apr  8 22:58:12 kkissling-laptop kernel: [ 25.979779] iwL3945 0000:03:00.0: firmware: requesting iwLwifi-3945-1.ucode
Apr  8 22:58:12 kkissling-laptop kernel: [ 26.081245] iwL3945: Radio disabled by HW RF Kill switch
Apr  8 23:03:26 kkissling-laptop kernel: [ 340.336109] usb 1-1: new high speed USB device using ehci_hcd and address 2
Apr  8 23:03:26 kkissling-laptop kernel: [ 340.612064] usb 1-1: configuration #1 chosen from 1 choice
Apr  8 23:03:27 kkissling-laptop kernel: [ 340.998382] Registered led device: rt73usb-phy1:radio
Apr  8 23:03:27 kkissling-laptop kernel: [ 340.998400] Registered led device: rt73usb-phy1:assoc
Apr  8 23:03:27 kkissling-laptop kernel: [ 340.998418] Registered led device: rt73usb-phy1:quality
Apr  8 23:03:27 kkissling-laptop kernel: [ 340.998860] usbcore: registered new interface driver rt73usb
Apr  8 23:03:31 kkissling-laptop kernel: [ 344.903790] rt73usb 1-1:1.0: firmware: requesting rt73.bin
Apr  8 23:03:31 kkissling-laptop kernel: [ 345.037774] ADDRCONF(NETDEV_UP): wlan1: link is not ready
kkissling@kkissling-laptop:~$
    
```

Figure 8: From the system log, you can discover the name of the driver for your WLAN card.

However, these drivers have not (yet) been integrated with the kernel, so you need to install them manually.

Unlike the official Linux kernel, Ubuntu supports WLAN cards and USB sticks with the RT2860 and RT2870 chipsets. If you have a device with an RT3070 chipset, you need to build the Linux drivers manually.

To build the drivers, start by using the Synaptic package manager to install the `linux-headers-kernelversion` and `build-essential` packages. (Replace `kernelversion` with your kernel version.) After downloading the drivers from the Ralink homepage [2], right-click the archive in the file manager to unpack in your home directory. The archive includes a `README_STA` file; the Build Instructions section describes how to modify the `config.mk` file in your favorite editor before compiling the software. In order for the driver to cooperate with the Network Manager, the `config.mk` file, which is found in the subdirectory `os/linux/config.mk`, must contain the following lines:

```

HAS_WPA_SUPPLICANT=y
HAS_NATIVE_WPA_SUPPLICANT=2
SUPPORT=y
    
```

After adding (or editing) these lines, change to the driver directory with `Makefile`, and type `make`. This command will build the module, which could take up to 30 minutes depending on your CPU. The `sudo make install` command then installs the module.

If this process completes without errors (warnings are okay and can be ignored), the driver should work. Copy the file `RT2870STA.dat` in the source directory to the `/etc/Wireless/RT2870STA` directory. Again using a text editor and armed with administrative privileges (`Alt+F2` and `gksu gedit`), you should now edit the `/etc/Wireless/RT2870STA/RT2870STA.dat` file to modify the transmission range for com-

pliance with your local requirements. (If you chose a different driver from the Ralink homepage, replace `RT2870STA` with the name associated with your driver.)

Take care of the following entries

```

CountryRegion=0
CountryRegionABand=2
CountryCode=US
WirelessMode=5
    
```

and see the `README_STA` file for a detailed description of the configuration.

Now reboot your machine and check to see whether the driver is working by first typing `iwconfig` in a terminal (Figure 9) and then typing `iwlist chan`. In the final section of `README_STA`, you will find out how to load the module automatically when you boot your machine.

If you have a USB stick with a chipset that the driver does not know, you will need to enter the USB ID in a file. First, launch an editor and open the `os/linux/usb_main_dev.c` file in the directory with the downloaded driver.

Now you need to scroll down to about line 100, where the list begins for the supported sticks and matching USB manufacturers and device IDs. This is where you enter your stick's ID. To discover your stick's ID, plug in the stick and type `lsusb` in a terminal. After completing these changes, build the driver by typing `make` and `sudo make install`.

Info

- [1] Ubuntu wiki: <https://help.ubuntu.com/community>
- [2] Ralink drivers for Draft N cards: <http://www.ralinktech.com/ralink/Home/Support/Linux.html>

Table 1 gives you an overview of well-supported USB sticks. Note that the manufacturers frequently change the hardware without renaming the models.

Legacy Burdens

Normally, Network Manager provides the easiest approach to setting up a WLAN device. However, Network Manager does not support certain legacy WLAN sticks and cards. If Network Manager is not cooperating, Ubuntu offers a manual configuration option. Start by clicking on the Network Manager icon and selecting `Edit Connections`. Now go to the `Wireless` tab, select the first entry, and click `Edit` to change the settings. In the new dialog box, first choose the `Wireless` tab, then uncheck `Connect automatically` and enter the required parameters.

Conclusions

Once you have taken the time to set up the adapter, opening a WLAN connection becomes easy. If you often work at different locations, Network Manager remembers individual connections and logs in automatically whenever possible. ■

```

File Edit View Terminal Tabs Help
carsten@systeme:~$ iwconfig
lo        no wireless extensions.

eth0     no wireless extensions.

wmaster0 no wireless extensions.

wlan0    IEEE 802.11bg  ESSID:"WLAN-00150CE0834E"
Mode:Managed  Frequency:2.462 GHz  Access Point: 00:15:0C:E0:83:4E
Bit Rate=54 Mb/s   Tx-Power=14 dBm
Retry  limit:7    RTS thr:off   Fragment thr=2352 B
Power Management:off
Link Quality=71/100  Signal level:-52 dBm
Rx invalid nwid:0  Rx invalid crypt:0  Rx invalid frag:0
Tx excessive retries:0 Invalid misc:0  Missed beacon:0

vboxnet0 no wireless extensions.

pan0     no wireless extensions.

carsten@systeme:~$
    
```

Figure 9: The `iwconfig` command displays information on wireless connections.

Manufacturer	Belkin	Buffalo	Conceptronic	D-Link	Hama	Hama	Hama	Linksys	Linksys	Sitecom
Model	F5D8053	WLI-UC-G300N	C54RU	DWA-140	39749	62778	62744	WUSB600N	WUSB54GC	WL-302 v1 001
Modes	b/g/n	b/g/n	b/g	b/g/n	b/g	b/g/n	b/g	a/b/g/n	b/g	b/g/n
Chipset	Ralink	Ralink	Ralink	Ralink	Ralink	Ralink	Ralink	Ralink	Ralink	Ralink
Driver	rt2870sta	rt2870sta	rt73	rt2870	rt73	rt2870sta	rt73	rt2870sta	rt73	rt2870sta
Interface	ra0	ra0	wlan0	ra0	wlan0	ra0	wlan0	ra0	wlan0	ra0
Price (approx.)	US\$ 55	US\$ 55	n.a.	US\$ 50	EUR 20	EUR 20	EUR 20	US\$ 65	US\$ 51	n.a.

n.a. means the device is no longer available.

Safety First

Linux is a very secure system, but it still pays to be cautious.

BY ACHIM LEITNER AND CARSTEN SCHNOBER

Security is one reason many users switch to Linux from Windows. With the rising importance of computers in daily life, more and more private data are stored on desktop systems. The hunt for bank account access keys has become big business, and many users also store private messages, business plans, and other sensitive information on home PCs.

A fresh installation of Ubuntu is much less vulnerable than a comparable Windows system. Spyware, trojans, viruses, and other attack vectors have not settled in the world of Linux as they have with other systems, but it is important to understand some basic security if you want to keep your Linux system safe.

Linux Security Basics

Every discussion of computer administration revolves around finding the balance between convenience and security. If you're careful, you can have a system that is both convenient and secure. The best approach is to let each user do what is needed, and nothing more.

This basic idea is most obvious in the multiuser system, which includes Linux and recent Windows and Mac OS systems. Every user stores documents, email, and other personal data in a private area. In Linux, this area is known as the user's *home folder*. If a user accidentally chooses something like *Delete all*, other users' data are not affected. The operating system itself is in a place where an individual user can't delete or edit the files.

Not only are misbehaving users dangerous, but malfunctioning software or an outside intruder who has breached security can cause considerable damage.

Root

Users can take care of their daily work without the need for write access to critical system files. Sooner or later, though, you'll need to make changes to the system. For instance, new software or hardware might require changes to the system configuration.

Every operating system has a special user account with the elevated privileges needed to change the system; in Linux, this special account is called *root*. The root account, which is also sometimes called the superuser account, corresponds to what Windows users call the *Administrator* account.

When you log into a Linux system as the root user, you are allowed to do virtually everything: You may read every document, de-

lete every file, remove every directory, and create new directories wherever you want.

Risky Business

Windows users who do their daily work as *Administrator* to avoid annoying privilege restrictions are taking a serious security risk. This warning also applies to the root user account in Linux. The root account should only be used for administrative work when it is required.


Even if you are charged with the task of administering the computer and are given the root password, you should still give yourself an ordinary user account for daily user tasks.

The traditional way to use the root account is to work as a normal user and, when necessary, switch to root during single administrative actions. Ubuntu further reduces the temptation to permanently work with superuser privileges by completely deactivating the root account by default. That means you will be able neither to log in as *root* nor to switch to the root account during your session.

Sudo

Of course, Ubuntu does not bar you from performing administrative tasks. Instead, the user account created during installation can gain superuser privileges with the command-line program *sudo*.

An advantage of *sudo* is that you do not have to remember another password. Before an administrator can execute an operation that requires privileged permissions, such as changing network settings, setting the time, installing



new software, or adding or removing software, Ubuntu asks for the user's normal password.

In one sense, this approach has the effect of reducing security for greater convenience because an intruder only needs to know your user password to gain full access to the system. Therefore, it is especially important that every person using the computer has a separate account. On the other hand, this approach forces the intruder to guess both a username and a password, whereas the classic root account always has the same name: *root*.

Managing Users

When you create a new account, you can choose whether the new user is a normal desktop user or an administrator who is qualified to perform system configuration tasks.

To create a new user, open the Users and Groups dialog from *System | Administration* to see a list of existing users. First, click the *Unlock* button and enter your password; if you have administrative rights, you may now add and delete users and manage groups.

If you choose *Add User*, a New User dialog (Figure 1) pops up. In the *Account* tab, enter

Figure 1: Adding a new user account.

Figure 2: Use the *Permissions* tab to set file permissions.

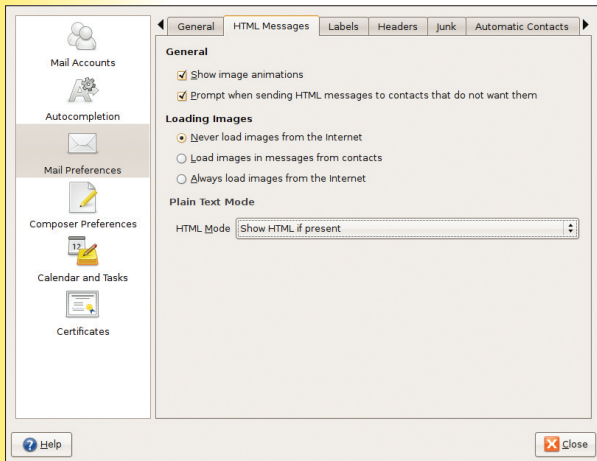


Figure 3: Make sure you know how your mail client is processing HTML email.

the *Username*, *Real name*, and optional *Contact Information* (e.g., telephone numbers).

In the *Password* section, you can either enter a temporary password yourself or make the system generate one. The new user is expected to change the password after the first login.

Granting Rights

In the *Profile* drop-down list, you can grant rights according to three levels of access. If you choose *Administrator*, the new user is allowed to gain superuser privileges and thereby can operate on the system without limits; otherwise, choose *Desktop user*. This lets the new user work in a home folder and gives full access to external media but forbids tasks that result in changes to the system settings. *Unprivileged* creates a user with minimal rights. A possible use for this kind of account would be to support users who log in via a network and therefore do not need access to local devices.

The *User Privileges* tab reveals more details about the chosen profile. Here, you'll find a list of specific actions, such as *Administer the system*, *Access local storage devices automatically*, or *Use audio devices*. All of these actions can be activated separately.

If you choose the *Administrator* profile, you'll see a check mark next to every entry. By clicking the checkboxes, you can adjust each new user's privileges individually. Be aware that users who are allowed to administer the system can change their own permissions.

Linux employs the concept of *groups* to manage access to files and other resources. In Unix and Linux, a group is a collection of users with certain common permissions. Linux grants different write and read permissions to the owner of the file and to the members of the group.

In the User Settings main window, the *Manage Groups* button shows a list of all groups in your system. Double-click an entry to see the group's current members. If you want, you can add users to the group manually, which has the same effect as activating the corresponding

The basic access categories are read, write, and execute.

When you right-click on any file in the file browser and choose *Properties*, the *Permissions* tab provides a view of access permissions (Figure 2). The upper section shows the file owner and the owner's access level: *Read and write*, *Read-only*, or *None*. Execution permissions are handled separately with the checkbox labeled *Allow executing file as program*. Whether you are able to change these settings depends on your access permissions.

Directories

Directories are handled like files except for execution. (Folders can have execution access as well, which means that users are allowed to browse the directory content.) The Properties dialog of a directory looks like that of a file, but with a button labeled *Apply permissions to enclosed files*. If you click this button, the permission settings affect files and subdirectories in this folder.

Filesystem settings can override file permissions. Ubuntu makes use of this feature for external storage devices. By default, external storage devices are mounted with the *noexec* flag, which means that enclosed programs cannot be executed. This protects the system from uncontrollable software brought in by CDs or other external sources.

In a terminal window, the command *ls -l* gives file permissions information, showing every entry in the current directory in a single line:

```
-rw-r--r-- 1 carsten 2
users 43142 2008-04-28 2
20:17 data.txt
```

checkbox in the user properties.

Read, Write, and Execute

User permissions are implemented as a system of access levels. This even applies to the access control on hardware devices because Linux manages these devices through special files residing in the */dev* directory.

File permissions are divided into three sections that specify what the owner, the members of the owner group, and all other users are allowed to do with the file.

For normal files, the first position is a dash; otherwise, *d* points to a directory. Device files and other special files show *c*, for example. The nine symbols that follow relate to permissions. These symbols are interpreted in three sections of three letters each. The first section refers to the file owner, the next to the group, and the last three to the remaining users. Each of these nine positions can show either a dash or one of the letters *r* (readable), *w* (writable), or *x* (executable).

Viruses in Linux

Theoretically, viruses could threaten Linux, but in practical terms, this hardly ever happens. The reasons are mainly rooted in the technology. To propagate, the virus needs to affect other programs it locates on the filesystem. Older Windows versions give viruses a free hand, whereas Linux slams the door shut. As long as the program has a normal user account, rather than one with root privileges, it is not permitted to overwrite other program files, leaving the virus without a propagation vector.

On top of this, Linux users rarely send each other pre-compiled binaries. Because any software you need is available for free, users are more likely to obtain the program from a website or distribution DVD.

Several vendors provide virus scanners for Linux [1], but the primary purpose of these tools is to protect email messages or files on file servers that might one day find their way to a Windows system.

Email Worms

In contrast to viruses, worms do not need a host program but run autonomously. A typical

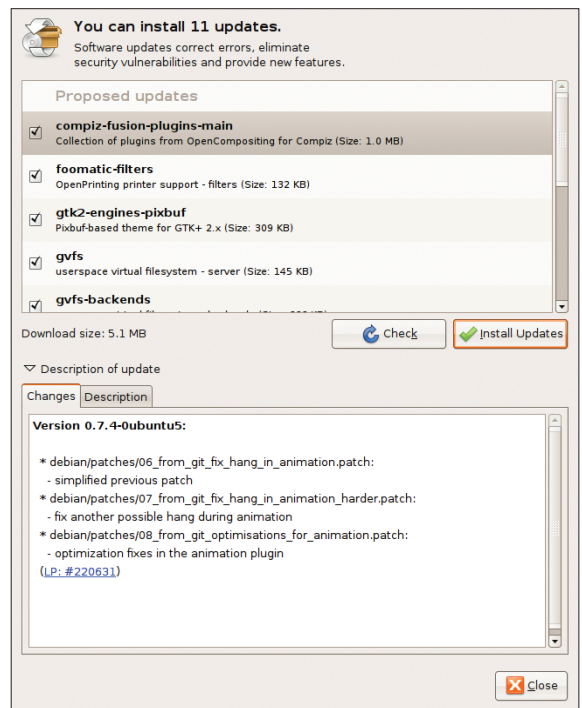


Figure 4: The Update Manager helps you keep your system up to date.

email worm sends itself to all the contacts in an address book. However, it only causes damage to the receiving machine if the recipient carelessly opens the attachment or if the worm can exploit a vulnerability in the mailer.

HTML-based email is also a risk. Many email clients ignore HTML attachments or deploy secure program modules to display HTML code. Be sure you know how your web browser treats HTML-based email. In Linux – just as in any other operating system – don't open attachments with untrusted messages and, above all, never launch any programs sent to you unexpectedly.

Evolution

Ubuntu's default mailer Evolution does display HTML mail by default. If you prefer to avoid this, use *Edit | Preferences*. In the *Mail Preferences* section, you will find the *HTML Messages* tab (Figure 3).

In the *Plain Text Mode* section, you can disable HTML by choosing *Only ever show PLAIN*. Some email messages consist of two

parts: One contains an HTML version, and the other contains pure text. If you prefer HTML, select *Show HTML if present*.

HTML email also provides the possibility of embedding direct links to Internet images. Because it is generally better not to load files from the Internet unless you know what they are, the default setting is *Never load images from the Internet*. If you indeed want to see images from a known secure source, change this option in the HTML preferences to *Load images in messages from contacts* or even to *Always load images from the Internet*.

Your own mail is written in plain text by default. For HTML, go to the *Composer Preferences* section in the Preferences window. To enable HTML formatting, activate the checkbox next to *Format messages in HTML*. Alternatively, you can enable it for a specific message in the composing window.

When you are writing a message, the Format menu provides the option *HTML*, which you can check or uncheck depending on current needs. Be aware that the receiver's mail software also has to be able to display HTML, so if in doubt, you should stay with plain text.

Don't Forget to Update

When a vulnerability is discovered, the correct response is to fix the program, preferably by updating the vulnerable program package. The Ubuntu Update Manager automatically locates patches for installed packages.

Vulnerabilities in server programs are particularly critical. The server just waits for a client to connect, and, of course, an attacker who gets inside can hitch up a server to the Internet and wait for victims. In fact, an attacker doesn't even need to know a victim's address: The intruder can simply search large sections of the network looking for low-hanging fruit. What makes this even worse is that servers often run with root privileges – following an attack, the attacker assumes the privileges of the victim, thus controlling the whole computer.

By default, Ubuntu checks for available updates every day. Ubuntu will indicate the presence of a new software package for your system with a message and put the Update Manager icon into the upper panel.

If you left-click the icon, the Update Manager window opens (Figure 4). Different update types might be available; when security updates are listed, you should install them.

Recommended updates usually fix bugs in applications that could cause application crashes but are not considered to be related to security.

After you click on the *Install Updates* button, the system asks for your password then downloads and installs the updates.

Firewalls

Many networks deploy firewalls to protect against outside attack (Figure 5). Firewalls only

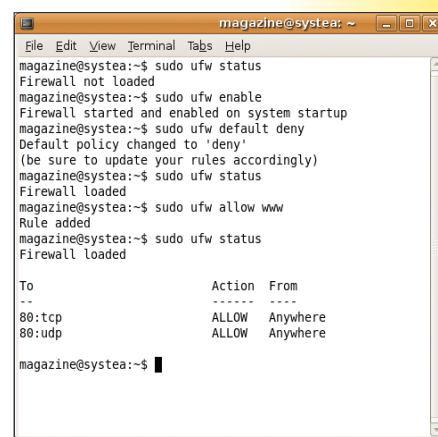


Figure 5: The ufw firewall configuration utility is useful if you are comfortable at the command line.

allow data packets to pass if they comply with specific criteria. A firewall can only give you limited protection: If you permit a connection, you are leaving it up to the client and server programs to defend themselves. In other words, the best protection is to install and launch only those services you really need.

In addition to server programs, which must be available to outside access and, therefore, are vulnerable to attack, some client programs are also endangered. A client application theoretically opens itself up to attacks by servers. For example, a web browser surfing the web exposes itself to the dangers of spoofing or phishing attacks.

As a protection against outside attack, Ubuntu installs no server software in the default configuration. However, many users might want to explore the possibility of running network services in Ubuntu.

If you are looking for some additional control, Ubuntu 9.04 comes with a firewall configuration tool called Uncomplicated Firewall (ufw). To configure ufw at the command line, enter the command *sudo ufw* followed by some arguments. The prepending of *sudo* is necessary because controlling network traffic requires administrator privileges. Complete documentation for ufw can be found on the Ubuntu wiki [2].

Solid Security

Linux does a good job of protecting itself against attackers, but you need to be careful and always install the updates. Although a little attention is warranted, the natural defenses of Linux are extremely successful at keeping intruders away. ■

Secure Passwords

Even Linux is powerless if users inadvertently or negligently endanger its security. Passwords play a key role in the security of a system. If you want a safe system, don't use passwords that are easy to guess. Using a friend's birthday or your pet's name is definitely not a good idea. An attacker who is looking to break passwords will deploy programs to automate the guesswork. These programs use comprehensive dictionaries in multiple languages and possess rules on how to compile words and replace letters with figures or nonstandard characters.

To make the guesswork as difficult as possible, a password should not be too short. Unix and Linux systems often demand eight-character passwords. These length restrictions are no longer valid for systems that use MD5-based password encryption techniques.

The number of different characters is particularly important: A mixture of uppercase and lowercase letters, numbers, and nonstandard characters, such as # < > _ = () , makes a password more effective.

Despite all the versatility and randomness, a user needs to be able to remember the password. Mnemonic aids can help, but they have to be just as secret as the password itself. Also of importance is that each password be used for one task only. For example, you should never use your login password for a web forum as your password for the root account on another computer. And you should change your password at regular intervals.

Info

[1] ClamAV open source antivirus tool: <http://www.clamav.org>

[2] Uncomplicated Firewall documentation: <https://wiki.ubuntu.com/UbuntuFirewall>

Printing

Before you set up a printer on Ubuntu, take a moment to learn about the Linux printing system. BY KRISTIAN KISSLING

If you have a fairly new printer, you can just plug it in, and it is likely to work. Before doing so, select *System | Administration | Printing*, which lists the printers the system supports. Then, connect your printer and check to see whether Ubuntu autodetects it and whether it appears in the list. If your computer is connected to a network and you are trying to find your office printer, click *New* and the system will check local connections, as well as your home or office network, for usable printers. (Make sure the printer is turned on.)

A summary of printers appears (Figure 1), hopefully containing the printer you are trying to configure. The five items at the bottom let you change the search parameters – depending on whether you are looking for an HP printer with driver support from the HPLIP project, an IPP printer that uses the Internet Printing Protocol, or an LPD/LPR server (Line Printer Daemon/Line Printer Remote) that uses TCP/IP to transfer data. On the right, you can enter the IP address of the computer to which the required printer is connected. If you use a network with static IPs, the printing system might be unable to find your printer.

If your printer is attached to a Windows computer on the network, you will need to en-



able file and printer sharing on the Windows computer and share the printer on the network. (Consult your Windows documentation.)

If the printer connected to the Windows system is properly shared on the network, it is typically sufficient to click *Windows Printer via SAMBA* in the *Devices* box and then click *Browse* to display a list of available Samba printers (Figure 2).

If the printer does not appear in the list, check the Linux Foundation [1] printer database for your printer. This well-maintained online database is a big help if you are looking

for the right driver, and it will tell you whether and how well Linux supports your printer.




Figure 2: Clicking the printer configuration menu is all it takes to find printers that Samba shares on the local network.

Theory

PostScript is a page description language originally invented by Adobe to provide a uniform set of instructions for all printers. PostScript Version 3 is the latest version, and it is mostly compatible with versions 1 and 3. If you have a PostScript printer, it will definitely run on Linux.

Before the advent of PostScript, every printer needed its own driver. Developers of graphics packages had to program, say, 20 printer drivers to support 20 different models.

PostScript solved this problem (Figure 3): Developers only needed to ensure that their applications created correct, readable Post-

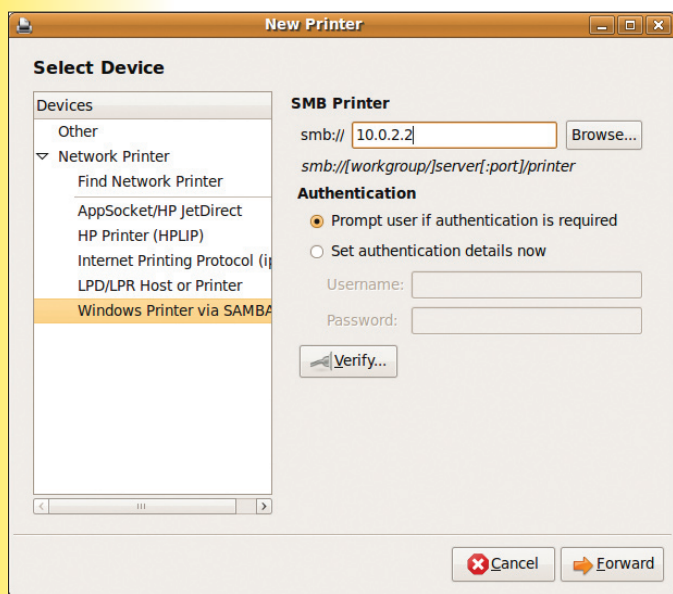


Figure 1: The Ubuntu printer dialog will automatically locate printers on your network or directly attached to your computer.

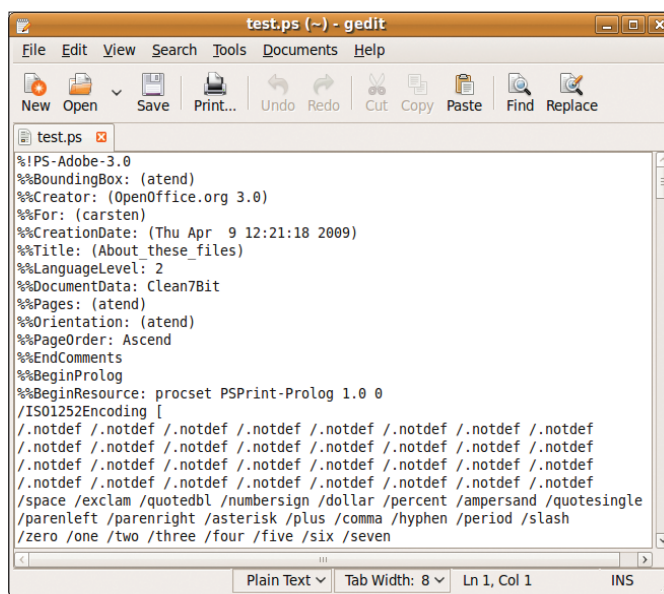


Figure 3: PostScript is a human-readable page description language for printers.

Script. The printer manufacturers did the rest by making sure their printers spoke PostScript. The Print Control Language (PCL) by Hewlett-Packard, the Graphical Device Interface (GDI) by Microsoft, and Epson's ESC/P were attempts by other manufacturers to develop alternatives to PostScript.

Before printing, a printer converts the PostScript input file into raster graphics – this procedure is known as Raster Image Processing (RIP). PostScript printers have their own intelligence when it comes to converting PostScript files into raster graphics. The cheaper models will have a cheaper chip on board, whereas the more expensive ones will have an embedded operating system that uses software to convert the input.

Printing Without PostScript

If your home printer does not support PostScript, the PostScript files created by your applications are sent to the conversion program first. The converter creates a raster graphics document that the attached printer can understand. The name of the Linux program that does this is GhostScript. GhostScript provides filters that have mainly been programmed in the course of a complex reverse engineering process, without any cooperation from the printer manufacturers.

This fairly complex process of analyzing the software structure makes it possible to install the software, even though the source code and documentation are missing. The results are not always perfect, but some filters – for example, the filters provided by the Gutenprint project [2] – now actually provide better quality than their Microsoft counterparts.

Enter CUPS

The Common Unix Printing System (CUPS) is regarded as *the* software tool for setting up printers on Linux (Figure 4). CUPS was originally based on GhostScript; however, at an early stage, it already had its own, highly functional filters. As of version 1.1.15, CUPS moved to ESP GhostScript 7.05, a commercial variant based on the GPL'd GhostScript. In August 2007, the two GhostScript projects joined forces to become GPL GhostScript 8.60, which now supports many non-PostScript printers on Linux.

The special thing about CUPS is that it can evaluate PPD files for printers. (See the box titled "PPD Files.") The descriptions are often integrated in Windows drivers and were originally only readable on Windows and Mac OS systems. Many Windows PPD files for PostScript printers will work natively on Linux thanks to CUPS, and CUPS even uses PPD files for non-PostScript printers.

Printer Resources

Despite the long list of working printers, certain printers still do not work with Linux. The

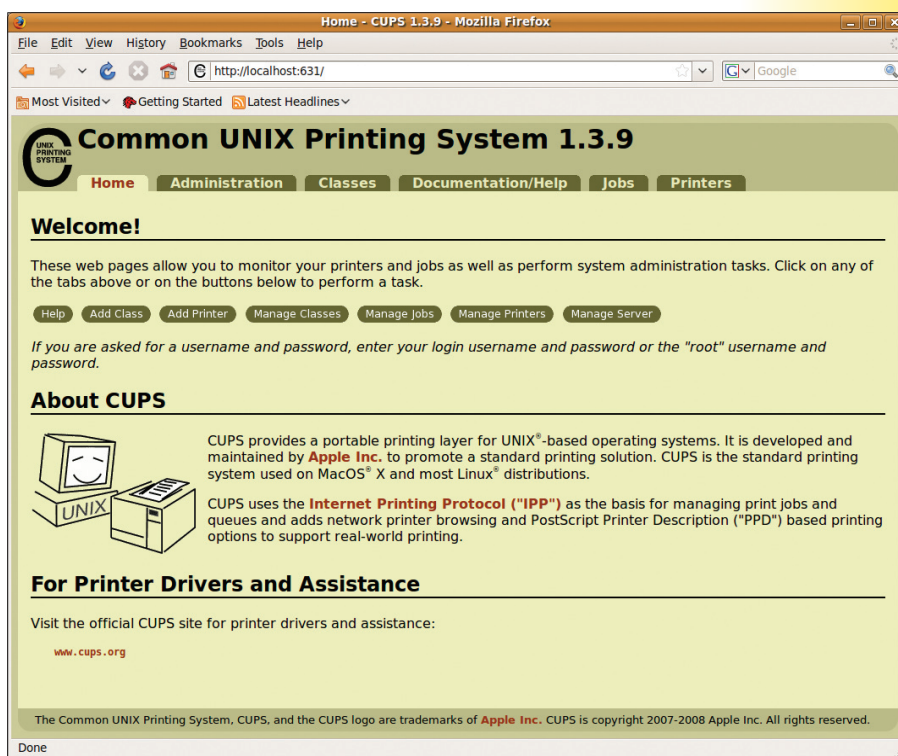


Figure 4: CUPS can be managed via a web interface by typing `http://localhost:631` in the browser address box.

number of printers on the market is just too big: So, if you're shopping for a printer, it's always a good idea to check whether or not the printer you are interested in has Linux support before you buy it.

The OpenPrinting.org Foomatic system [3] lets you convert legacy GhostScript filters into PPD files. The use of this approach will not necessarily improve the printing quality, but CUPS offers the benefit of easier printer management, and it works better with printer management GUIs.

To use the Foomatic filters, you need to install the packages on Ubuntu with *foomatic* in their names. As a general rule, before you start building from the source code, you need to check first to see whether you can find a CUPS package with the required PPD files for your printer.

If you are not satisfied with the print quality, it is a good idea to check OpenPrinting.org to see whether an alternative PPD file for your printer exists. The website should be your first stop if you are shopping for information on

PPD Files

The printer driver evaluates the PostScript Printer Description (PPD) file to discover the printer's capabilities. (For example, the file describes whether the printer supports features such as duplex printing.) Printer management GUIs also reference these files to glean information on what buttons, widgets, and options to show the user.

printing on Linux. It has a huge collection of documents on the subject – and the CUPS project itself also has information pages [4]. ESP PrintPro is the name of the now-defunct commercial variant of CUPS that supported thousands of printers. Work on this tool was discontinued when its main developer, Michael Sweet, moved to Apple.

If you search the Linux Foundation pages for your printer, you might be taken to one of the numerous minor sites that offer printer drivers. For example, the Gimp-Print project, in collaboration with Gutenprint [1], supports around 700 printers, including HP Laserjets and Deskjets, Epson Stylus and Photo Color models, and some Canon and Lexmark devices. In addition, Ubuntu has its own Gutenprint package (*foomatic-db-gutenprint*) with the drivers needed by CUPS. ■

Info

- [1] Finding printer drivers: http://openprinting.org/printer_list.cgi
- [2] The Gutenprint/Gimp-Print project: <http://gimp-print.sourceforge.net>
- [3] Foomatic project: <http://www.linuxfoundation.org/en/OpenPrinting/Database/CUPSDocumentation>
- [4] Background for PostScript, GhostScript, and CUPS: <http://www.linuxprinting.org/kpfeifle/LinuxKongress2002/Tutorial/III.PostScript-and-PPDs/III.PostScript-and-PPDs.html>

Multimedia Codecs

To experience the full spectrum of multimedia file formats available online, you might need to install some codecs. BY TIM SCHÜRMANN



Ubuntu offers a range of options for playing and creating multimedia files. The Rhythmbox music player (Figure 1) is installed with the default desktop. To play and organize your music collection with Rhythmbox, choose *Applications | Sound & Video | Rhythmbox Music Player* in the Ubuntu main window.

Totem is Ubuntu's default video player. To reach Totem, select *Applications | Sound & Video | Movie Player*. See the respective Help menus for more on how to get started with Rhythmbox and Totem. Both apps come with friendly and useful documentation.

Other multimedia tools are available from the Ubuntu repositories. See the Ubuntu wiki for more on the multimedia tools available for Ubuntu, or click the *Sections* button in the Synaptic package manager and scroll to the *Multimedia*, *Multimedia (multiverse)*, and *Multimedia (universe)* package sections.

For both legal and practical reasons, the number of audiovisual file formats supported in Ubuntu's default configuration is quite limited. If the file you want to play isn't supported, you might need to add a codec for the file format. In this article, I look at how Ubuntu handles codecs for multimedia.

In the Codec Jungle

When you see a video file on your desktop, you can normally launch the matching player by double-clicking. On Ubuntu 9.04, Totem is the application of choice for this job – but don't be surprised if this doesn't work on a system that you have just installed from scratch. By default, Ubuntu only plays free file formats, such as the MP3 competitor Ogg/Vorbis. Exotic

file formats from the world of Windows are likely to cause problems.

The best way to understand this problem is to consider that the program known as Totem is actually just a graphical user interface for the similarly modular GStreamer multimedia software. GStreamer in turn references plugins to learn about various file formats. Each plugin contains what is known as a codec. As the acronym for COder / DECoder suggests, a codec is a piece of software that converts audio or video data from one format to another.

For example, if Totem wants to play an MP3 file, it passes the data on to GStreamer, which checks its repository for a plugin that can handle MP3s. With this codec's plugin, GStreamer converts the file to audible material, which Totem picks up and outputs via the loudspeakers. To teach Totem new file formats, you simply need to install matching plugins (Figure 2). The question is: Which ones?

Onion Skins

File formats from the world of Windows, typically with suffixes such as *.avi*, *.wmv*, or *.asf*, are really just containers for the video and audio data. In fact, you can't tell whether the file contains an MPEG4 video or a file in the defunct Indeo format. Incidentally, the same thing applies to QuickTime files (*.mov* suffix), as produced by Apple computers.

When Totem finds an unrecognized file format, it will usually search for a matching GStreamer plugin before installing the corresponding Ubuntu packages (Figure 3). Totem and Rhythmbox

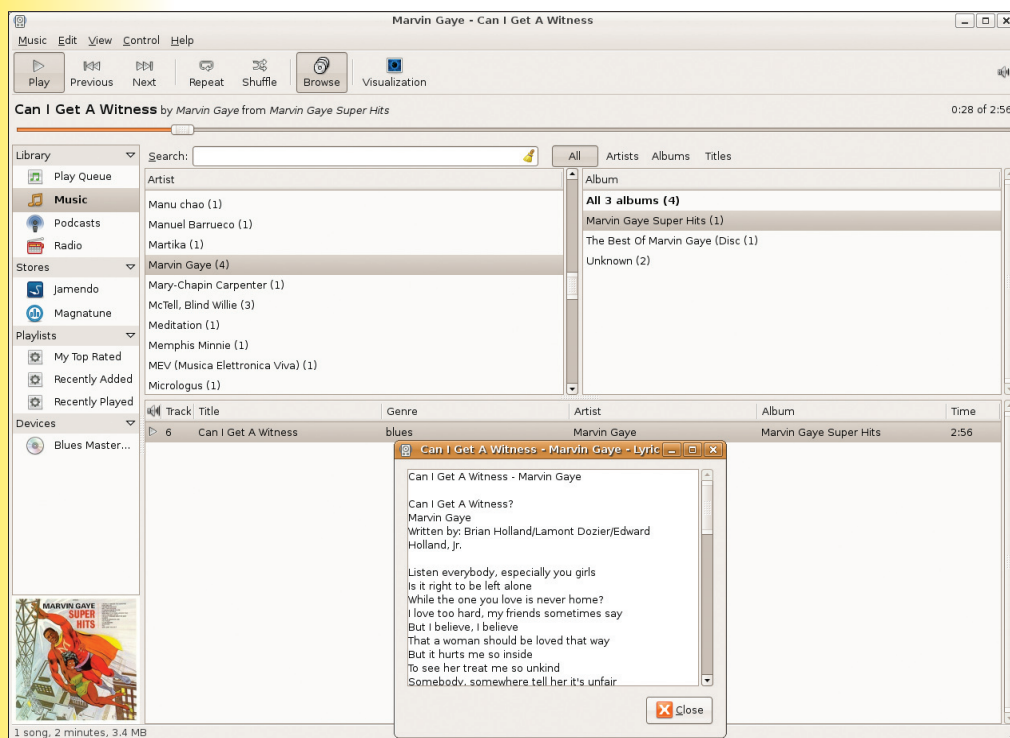


Figure 1: Manage your sounds with the Rhythmbox music player.

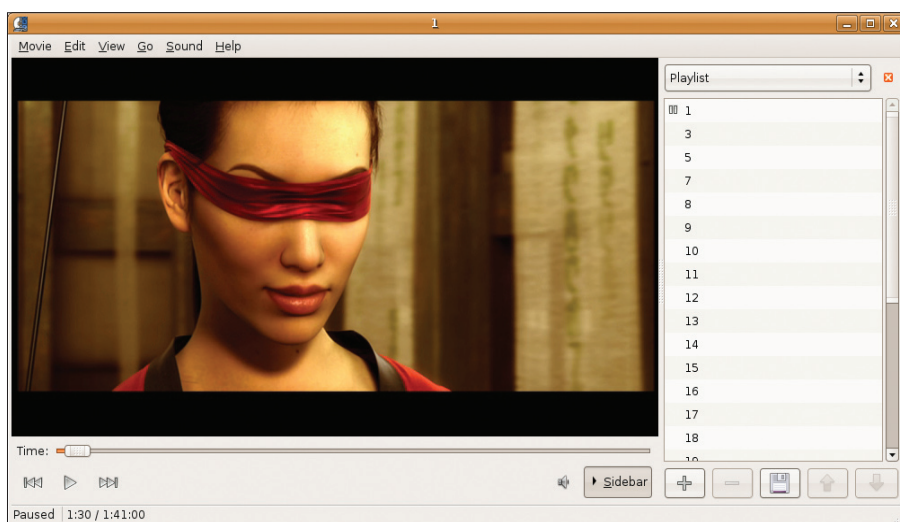


Figure 2: Play videos with Totem.

(which also relies on GStreamer) need to install the packages *gststreamer0.10-plugins-ugly* and *gststreamer0.10-fimpeg* to play MP3 files.

To avoid these questions in the future, you can apply a liberal dose of plugins by telling Synaptic to install the *ubuntu-restricted-extras* package. (This helps Ubuntu fetch numerous, frequently needed plugins for tools, such as Sun Java, or various Windows fonts, which can be quite bulky.)

A package called *gststreamer0.10-plugins-ugly* includes codecs for files with some legal restrictions. Many file formats are subject to patents and licenses (for more information on this issue, see the “Dangerous Liaisons” box). To identify plugins with a dubious legal status, the developers called them just plain *ugly*. If Totem offers to install an “ugly” package, you have to decide whether to install it and thus enter the legal gray zone.

Multimedia with Medibuntu

Some exotic file formats from the world of Windows remain closed to Ubuntu even if you

follow the steps outlined just now. In that case, you can retrieve the required codecs from an existing Windows installation, although this can be tantamount to the labors of Sisyphus. Fortunately, the external Medibuntu repository offers a package that bundles the major Windows codecs. To load the package, go to *System | Administration | Software Sources*, click the *Third-party software* tab, and then click *Add*. Now type the following, slightly cryptic line in the input box

```
deb http://packages.
medibuntu.org/ jaunty
free non-free
```

to add the software package source (pay attention to the spaces in this command line). When you leave the window by clicking *Close*, you are taken to the main window, where you can click on *Reload*. This action tells Ubuntu to pick up a list of available Medibuntu packages off the Internet. If you are warned about missing signatures, just click *Close* and try again a couple of minutes later.

When you are done, you need to install the *w32codecs* package.

The Pitfdll Dilemma

Sad to say, this is not the end of the story: You also need Totem support for Windows codecs. The Pitfdll plugin was developed for GStreamer as a kind of adapter that runs the Windows codecs on Linux. Synaptic calls the package *gststreamer0.10-pitfdll*. Unfortunately, installing it will not get you very far: The gang of three – Totem, GStreamer, and

Dangerous Liaisons

If you want to view commercial DVDs on Linux, you currently need a special library (*libdvdcss2*) which removes the CSS (Content Scrambling System) protection. In many countries of the world, the use of copy protection removal systems is illegal and could constitute a copyright offense. The only legal way to play CSS-protected disks right now is LinDVD by Corel [1], which costs around EUR 40.

Besides the copy protection issues, there are a number of unresolved patenting issues, relating first and foremost to compression methods such as MPEG2, or MP3, a format popular with music lovers. The players and libraries were developed without the explicit approval of the patent and/or license owners, and they are thus on uncertain ground from a legal point of view. Therefore, you have to install the packages manually – and at your own risk.

Pitfdll – do not work properly on Ubuntu 9.04 as of this writing.

To work around the problem, you can install an alternative version of Totem that uses the Xine media player library instead of GStreamer. To do so, launch Synaptic, search for the *totem-xine* package, and install the package. Now you have two versions of Totem on your disk.

The main menu will still take you to the GStreamer version, so to launch a movie with the Xine version, go to *Applications | Accessories | Terminal*, type *totem-xine&*, and press Enter. The video player does not look any different, but if you check *Help | Xine*, you will see that a different engine is under the hood. To play a video file directly with Totem Xine, right-click the file, select *Open with other application*, and then select *Video Player (Xine)* from the list. ■

Fluendo GStreamer Plugins

Fluendo [2] has licensed GStreamer plugins for sale on its website. In addition to the MP3 codec, you will find everything you need to play Windows media files (WMV, WMA, ASF) and MPEG2 videos, although a full set of codecs costs EUR 28.

Info

- [1] LinDVD player:
http://www.intervideo.com/jsp/Product_Profile.jsp?p=LinDVD
- [2] Commercial GStreamer plugins:
<http://www.fluendo.com/>

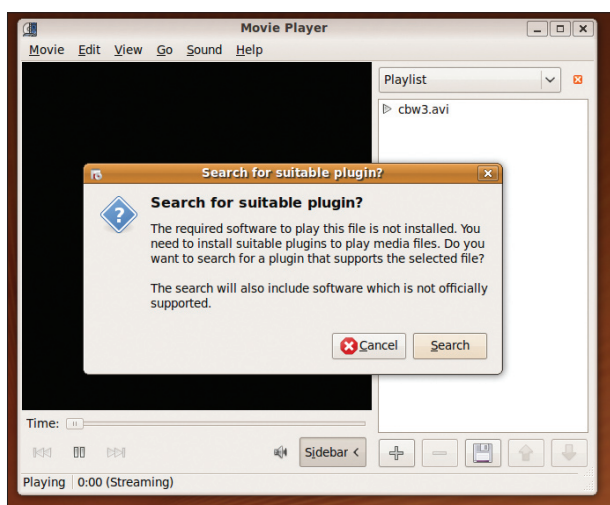


Figure 3: After a short pause for thought, GStreamer-based programs like Totem and Rhythmbox offer to install any missing codecs.

Burning CDs and DVDs

Burn your own CDs and DVDs with Brasero. BY MATTHEW SACKS

Ubuntu includes the latest release of Gnome's CD/DVD burning tool, Brasero [1]. This application is similar to Roxio's disc burning suite, but it is simpler and much easier to use.

To access Brasero, click on *Applications | Sound & Video | Brasero Disc Burner* (Figure 1). The main Brasero project window appears with the following project types: audio discs, disc copy, data disc creation, and disc image burning.

Burning MP3s

Brasero allows the user to drag and drop MP3s and most other audio file formats into a project to be burned to an audio CD for play in any standard CD/DVD player. The quality of the blank CD/DVDs you use can have a great effect on the system's ability to make a good burn – so try to avoid the discount blanks.

To begin creating an audio CD, click on the *Audio Project* button in the main Brasero menu or click on *Project | New Project | New Audio project*. Now drag the desired MP3 (or other supported audio format) to the main window or click the *Add* button and locate the audio files on the local hard drive (see Figure 2). At the bottom of the application you will see how much space the added files consume.

In the *Name* text field, you need to enter the title before you click *Burn*. A Disc Burning Setup dialog box appears, prompting you to use the burning device. To begin burning the audio project, click *Burn* again. A status bar

shows the progress of the burn as the audio files write to disc.

Disc Copying

Brasero has a disc copy feature that is useful when exact copies of the same disc are needed, whether it is a data disc, a DVD movie, or an audio CD. Because Brasero stores a temporary image of your CD or DVD on the local hard disk and then prompts you to insert a blank CD, just one drive is necessary.

To copy a disc, insert the source disc and have a blank CD/DVD handy. After clicking on *Disc Copy*, select a source drive to copy from and a destination drive to write to (Figure 3), then click the *Copy* button. A status bar indicates its progress. The program opens the disc tray and displays a dialog box prompting the user to insert a blank CD or DVD. The Waiting

for Disc Insertion dialog disappears once a valid blank disc is inserted into the drive tray and the tray is closed. As soon as the disc is finished copying, a message will display *CD successfully copied*. To view the contents of the CD or DVD, you need to remount it by simply opening and closing the drive tray.

To make copies of DVD movies (legal copies, of course), you need libdvdcss2 [2]. This library can be installed in Ubuntu by opening a terminal window and typing `sudo /usr/share/doc/libdvdread3/install-css.sh`. Totem, Ubuntu's default video player, also requires the libdvdcss2 library to play DVD movies.

Burning a Disc Image from an ISO

Brasero supports burning from an ISO image [3] to a CD or DVD. Disc images are a conve-



Figure 1: Starting the Brasero application.

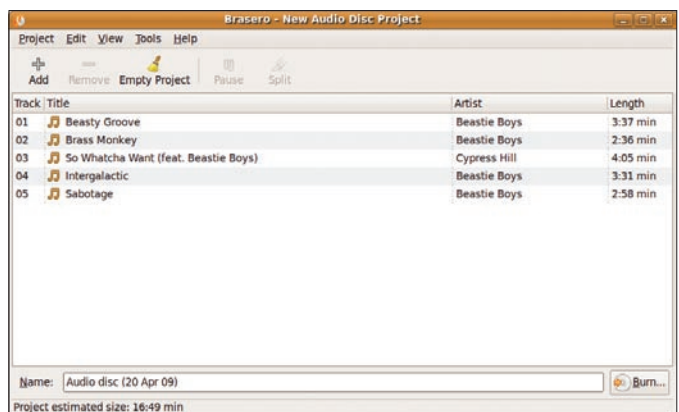


Figure 2: Adding music to an audio disc project.



Figure 3: The CD/DVD Copy Options dialog.

nient way to transport the contents of CDs and DVDs acquired from the Internet. For example, you could download the latest copy of a new operating system, such as Ubuntu, and burn it to a disc, or you could transfer audio files as an image and burn a copy to a physical disc for use in a DVD player.

In the main Brasero menu, click on the *Burn Image* button or click *Project | New Project | Burn Image*. The Image Burning Setup dialog box will appear (Figure 4).



Figure 4: Burning a disc image.

Now select the drive to use for writing the image to disc and click on *Click here to select an image* to locate the image on the local hard drive. The Select Image File dialog box will appear. Here, select the ISO image to be burned to disc, click *Open*, then click *Burn*. Brasero creates an image checksum and shows the progress of the image as it is burned to disc. If Brasero throws errors, try lowering the burn speed and then re-burning.

Creating a Data Disc

Brasero easily allows the user to create data discs to back up and transport files. Creating a data disc in Brasero is simple. Click on the *New Data Project* button or click on *Project | New Project | New Data Project*. Now add your documents to the *Files* pane and click the *Burn* button in the bottom right-hand corner. The Disc Burning Setup dialog box will appear. To change properties, like the burning

Info

- [1] Brasero: <http://www.gnome.org/projects/brasero/>
- [2] libdvdcss2: <http://en.wikipedia.org/wiki/Libdvdcss>
- [3] ISO image: http://en.wikipedia.org/wiki/ISO_disk_image

speed, use the *Properties* button (Figure 5). When finished, click *OK* and then *Burn* again. The progress bar will show the status of the burn. Unless you have changed the corresponding option, the disc should automatically eject from the tray once the burn is complete.

Conclusion

Brasero makes the process of burning CDs and DVDs extremely simple. The plugin capability of Brasero offers a means for integrating additional features. ■

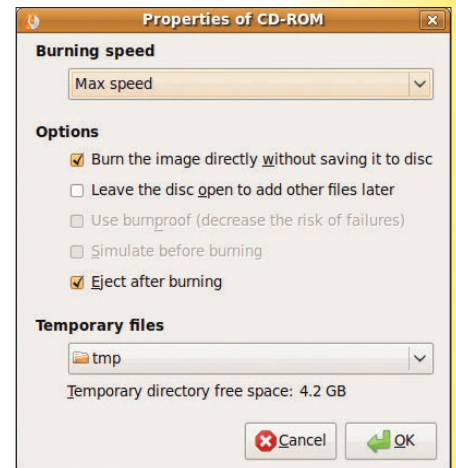


Figure 5: In the Properties of CD-ROM box, define the burning speed and other options.

DON'T MISS A SINGLE ISSUE!

The first print magazine created specifically for Ubuntu users! Ease into Ubuntu with the helpful Discovery Guide, or advance your skills with in-depth technical articles, HOW-TOs, reviews, tutorials, and much, much more.

SUBSCRIBE NOW!

4 issues per year for only
£ 24.90 / EUR 29.90 / US\$ 39.95

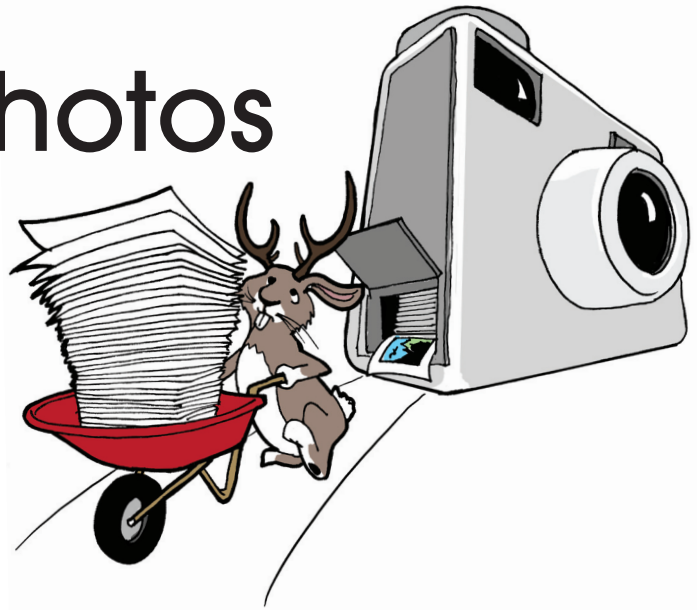
- ✓ Don't miss a single issue!
- ✓ Huge savings – Save more than 35% off the cover price!
- ✓ Free DVD – Each issue includes a Free DVD!



www.ubuntu-user.com

Managing Photos

Chaos is inevitable when your hard disk fills up with digital photos. The F-Spot image management program helps you manage your photo archives. BY MATTHIAS WARKUS



With the recent popularity of digital photography, hard disks tend to get swamped with photos. The problem with storing images on a computer is finding the photo you need in a reasonably short time.

F-Spot [1] can bring order to your personal photo collection. The project was launched in 2003 by Gnome developer Ettore Perazzoli, who unfortunately passed away in December of that same year. Now, Larry Ewing (who designed Tux the penguin) has taken over.

In recent Ubuntu versions, F-Spot is part of the default installation. Because the software is written in C#, it requires the Mono environment [2] and the matching program interpreter, *mono*, along with a few Gnome packages, various databases, and some graphics libraries.

Because Mono has become a part of the Gnome desktop, these packages would be included in a default Ubuntu installation anyway, so integrating F-Spot was a logical next step.

Getting Started

To start F-Spot, click on the Ubuntu Applications menu and choose *Graphics | F-Spot Photo Manager*.

When first launched, F-Spot displays a completely empty image catalog. However, the program does pop up a dialog from which you can import new photos. Then you can select a drive, an attached storage medium, or any folder as your import source. At the same time, you can define categories for the imported images. When you click *OK*, F-Spot will search the specified folder and catalog the image files it contains. Any time later, you can open the import window from the File menu.

Search and Find

After importing a few images, F-Spot should be slightly more colorful. When you click an image, details such as the name, size, or date are displayed (Figure 1).

F-Spot retrieves data, such as the exposure time or orientation, from the image files through Exif, a standard for transferring image-specific metadata. The Exif fields are displayed

in the Metadata Browser dialog box, which is in the *View* menu.

The jumble of images can be organized in two ways. First, you can use the tags in the list shown at the top left of the F-Spot window. Besides the special tags *Favorites* and *Hidden*, three categories are predefined: *People*, *Places*, and *Events*.

To specify a new category or to add new tags, you can use either the context menu or the *Tags* entry. Dragging a tag or category to the thumbnail assigns the image to the tag. Also, you can use the context menu for tag management by selecting either *Add tag* or *Remove tag*.

When you enable a checkbox in the category list, F-Spot displays the images from that category in the index. Enabling multiple checkboxes tells F-Spot to display images with *one* of the tags, but not the intersection of the selections. Selecting *Edit tag* in the context menu (or *Tags | Edit selected tag*) pops up a window in which you can change the category assignment of a tag.

The second organization method, a slide control, gives quicker access to images without requiring the user to tag the images first (Figure 2). This tool, which is quite puzzling at first sight, shows a time index with markings for years and months. To scroll the visible segment, click on the arrows.

A bar above the segment for a specific month tells you that images from that month exist in the current catalog, and the number of images is reflected by the size of the bar. If you

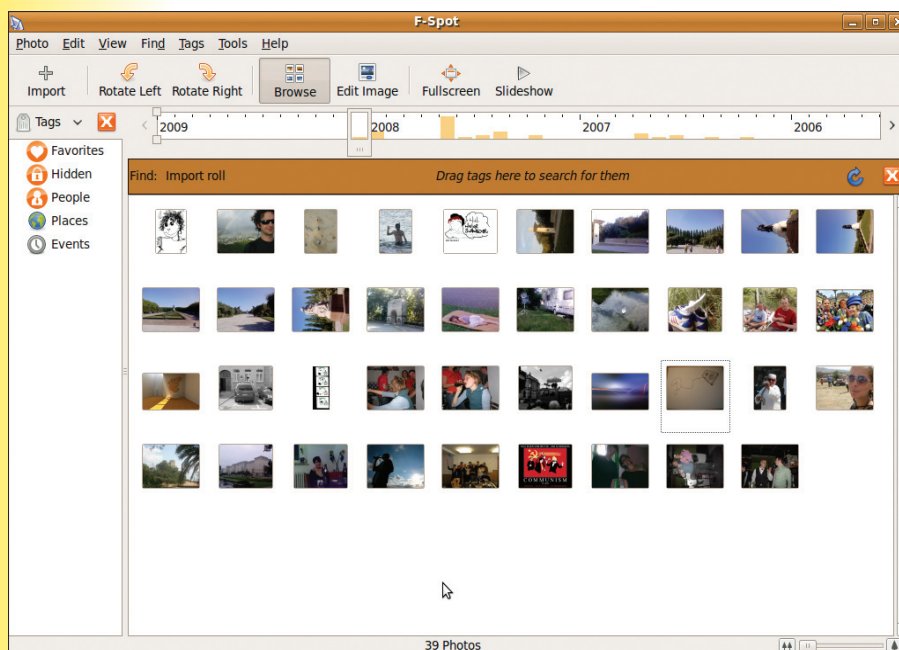


Figure 1: F-Spot shows an overview of stored images.

Info

[1] F-Spot homepage: <http://f-spot.org/>

[2] Mono download: <http://www.mono-project.com/Downloads>

[3] Flickr: <http://www.flickr.com/>



Figure 2: Time sliding.

then push the slider to a specific point in time, the index jumps to the image closest to that point. Also, you can use the handles to move the lower and upper date thresholds so that you restrict the index to a specific period.

If you select the folder view (*View | Arrange by | Folder*), F-Spot changes the timescale into a directory scale. This option can be very useful if you sort your images into folders before importing them.

Getting There?

After you have found the image you are looking for, you need to double-click it to display it in a scrollable, zoomed view (Figure 3). The buttons for typical features – such as rotate left, rotate right, or flip portrait images to landscape images – are available in the index view. In zoom mode, you can additionally use the frame tool to select a detail or remove red eye from portraits.

Also present is a list of aspect ratios that allows you to crop images to fit a specific page format. The palette is rounded out by a color histogram, a soft-focus tool, color adjustment, sepia tones, black and white conversion, and

automatic horizon straightening. Also, you can use the powerful sharpening tool by pressing *S* or selecting *Edit | Sharpen*.

When you choose to edit an image, F-Spot always creates a new version of it to prevent damage to the original. If things go wrong, you can select *Original* in the Version menu to restore the original image.

The program also has tools that you can use to select the current image as wallpaper. A full-screen view mode and a slideshow feature are also available. A printing function that uses the Gnome-print API to create hard copies gives users the ability to print their images if needed.

The Export menu provides more functions for sorting and editing images. Besides copying images to a different directory, it also supports

the Flickr [3] photo sharing service, Picasa Web Albums, SmugMug, and 23HQ, as well as allowing users to create and upload HTML galleries.

Users can burn their image collections onto CDs or DVDs with the Gnome burning interface.

Conclusions

F-Spot might be just what you need to bring order to the chaos of your photo collection. If your collection has grown over the years, you will like the timescale feature. ■




Figure 3: Zoomed view.

Freedom:

the freedom to run the software as you wish.

Join the FSF as an associate member.

Your donation will support our 20-year-old commitment to community.



Installing Software

Ubuntu lets you install and manage software at the click of a mouse. BY JAMES MOHR

Thousands of open source software applications are available for Linux systems, and because these tools are mostly free, the software usually isn't distributed in shrink-wrapped boxes. A Linux user who wants a program simply downloads it. In the old days, this meant finding and compiling the source code. Some Linux users still like to work with source code, and in some cases, your only choice might be to compile the program yourself. However, most users install software with package systems. A package system is a collection of tools and services that lets you download and install an application in just a few mouse clicks.

Major Linux vendors make packages available in large online repositories. A program on the client computer simply calls for a package in the repository, and the installation occurs nearly automatically. Best of all, the package system also resolves dependencies associated with the application. A dependency is a background application or component that must be present for the program to run. For instance, an application might require the correct version of a specific Linux library. A good package system installs all the necessary background software along with the application you want.

Ubuntu and other Debian derivatives use the Advanced Packaging Tool (APT) system. The Ubuntu project has a large repository of packages you can install on your Ubuntu computer. Several tools can serve as a front end for accessing these APT-based repositories. From a graphic desktop, Ubuntu provides the Synaptic package manager. Working from the command line, you can install a new package in one or

two commands with a utility called *apt-get*.

In this article, I look briefly at managing packages with Synaptic and *apt-get*, but you'll have to decide for yourself which applications to install. The economies of open source don't leave room for a lot of vendor advertising. Most Linux users find interesting tools by word of mouth. You might learn of a promising application in an online forum or from a magazine like this. If the app is available as an Ubuntu package, a user can quickly add it to the system with the use of Ubuntu's package management tools.

Installing Software

The Synaptic package manager is a graphic tool for installing, updating, and managing software packages. To start, select *System | Administration | Synaptic Package Manager*, then enter your administrator password.

The Synaptic main window (Figure 1) is where you tell Synaptic to display packages by subject, status (installed or not), or origin (which repository). Also, you can define your search criteria. As shown, I have already clicked the *Sections* button to display the packages in predefined groups. In this case, I have selected the *Science* section because I want to install the KStars planetarium program.

After I select the KStars package, a new window pops up showing a list of changes the system needs to make for the program to run (Figure 2). This list displays the dependencies associated with the program; if you want to remove a particular package, you might see a list of packages you need to *remove* to keep the system consistent. If you click *Mark*, Synaptic will mark all of the necessary packages for installation.

Clicking *Apply* in the toolbar at the top of the Synaptic main window pops up a Summary box showing all of the packages Synaptic is to install (Figure 3).



Often, some of the package have to be downloaded. Synaptic can find the necessary packages in different repositories and servers on the Internet. When you click *Apply*, Synaptic downloads the files and performs the installation, including creating directories and copying the files to the appropriate locations.

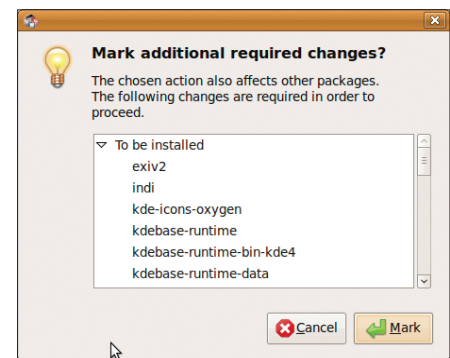


Figure 2: Synaptic summarizes the changes necessary for installing an application.

The time needed for installation depends on the number of packages, the speed of the Internet connection, and the extent of the necessary preparations. In this example, it took just a couple of minutes. When it's done, KStars appears under *Applications | Education*.

Keep in mind that Synaptic displays not only direct dependencies, but also the packages on which these dependencies depend.

Occasionally, you might not know the package name or even which group the package is in. Synaptic includes a search feature that helps you locate a package if you don't know the name. For example, you can search for a string of characters that appear in the package name or search by dependency or maintainer (the person who manages the package).

Other Features

Synaptic provides a number of other interesting features. For example, you can examine the details of each package by right-clicking on it

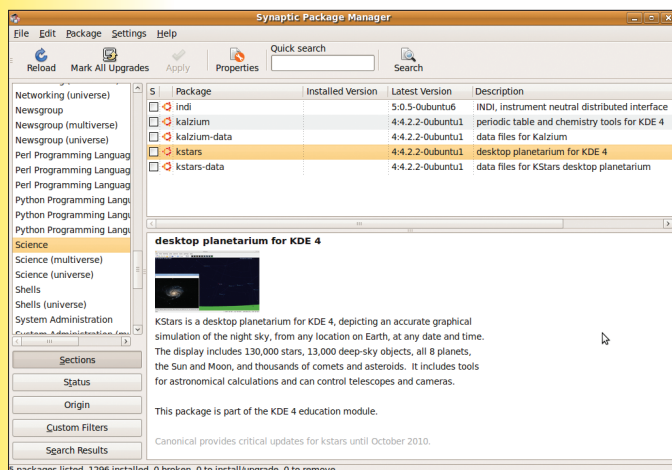


Figure 1: The Synaptic main window.

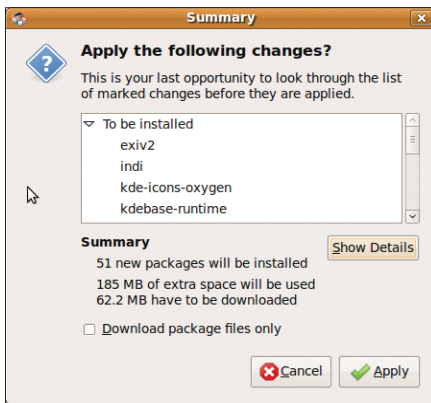


Figure 3: Study the list of changes before you click *Apply* to begin the installation.

and selecting *Properties* (Figure 4), which shows several useful settings. The *Dependencies* tab is useful for troubleshooting failed installations. Note that the final two entries in Figure 4 are labeled *Suggests*. These packages are recommended but not required for the program to run successfully. One common example of a suggested package is documentation for the program.

Synaptic also lets you track the history of what has been installed. To view the history,

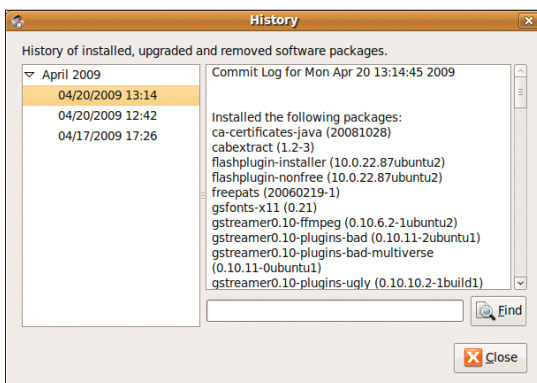


Figure 5: A view of the installation history.

select *File | History*, then select a specific entry, as shown in Figure 5.

Adding Repositories

After you finish installing Ubuntu, the system adds the installation CD to the list of repositories, along with the official repository Ubuntu maintains on the Internet. Under the *Edit* menu is an entry called *Add CD-ROM*, which adds a CD-ROM to the system. When you select this option, you are prompted to insert the CD, which is then scanned. The packages Synaptic finds are added to the list of known packages. Later, when you attempt to install a package that is present on the CD, Synaptic will prompt you to insert the CD.

The official Ubuntu repository is divided into four components:

- *Main* – free software supported by Ubuntu.
- *Restricted* – Ubuntu-supported software that is not available under a free license.

- *Universe* – community-maintained software that is not officially supported.

- *Multiverse* – unsupported software under a non-free license.

Synaptic also lets you add repositories located on the Internet or on a local network. In *Settings | Repositories*, you have a wide range of options for configuring repositories. For example, you can download only officially supported software or only community software, or you can add personal repositories that others can't access.

At the Command Line

Sometimes users need to install software directly from the command line, such as when connecting remotely or when using a system without a graphic desktop. Some system administrators prefer to work at the command line because they find it faster and more intuitive. Ubuntu uses *apt-get* to install software at the terminal prompt.

To install KStars from the command line, you must invoke the *sudo* utility, which is used with commands that require administrative privileges. In a terminal session, enter:

```
sudo apt-get install kstars
```

Note that *kstars* needs to be lowercase. The output in Figure 6 is similar to what you should see.

To initiate the installation, press Enter. *Apt-get* runs through the same steps as Synaptic. Installing software with *apt-get* is usually very easy. As you might guess, you can also use *apt-get* to perform other tasks, such as uninstalling packages and displaying dependencies. For more on working with *apt-get*, type *man apt-get* to view the *apt-get* man page.

Manually

Sometimes the package you want to install is not available in a repository, so you need to take an extra step to install the package “manually.” Ubuntu uses the Debian package format. Debian packages typically have a *.deb* ending. If you want to install a Debian package, simply double-click on the file in the Nautilus file manager, and a window will pop up allowing you to install the package. Your system will attempt to resolve any dependencies through your configured repositories. If the supporting packages are not available, you will need to search for the missing packages, which could be a tedious task.

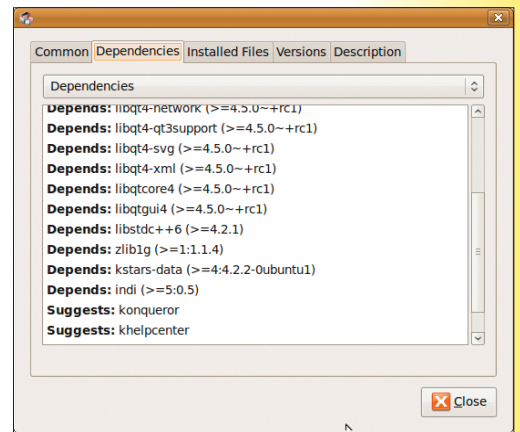


Figure 4: Synaptic displays package dependencies in the *Properties* dialog.

Debian packages also can be installed directly from the command line with the *dpkg* command:

```
dpkg -i PACKAGE_NAME.deb
```

One disadvantage of this approach is that you need to resolve the dependencies yourself, but this technique is sometimes useful if you are installing software remotely.

If a native *.deb* package is not available, you might find the program in an alternative Linux package format, such as an RPM. If you want to install a package in a foreign or “alien” format, try the *alien* command:

```
sudo alien -i PACKAGE_NAME.rpm
```

If the *alien* utility is not present on your system, find and install it using the techniques described in this article.

Conclusion

To familiarize yourself with the applications available for Ubuntu, browse the Ubuntu repository – and keep surfing and reading to learn about new and interesting Linux tools. ■

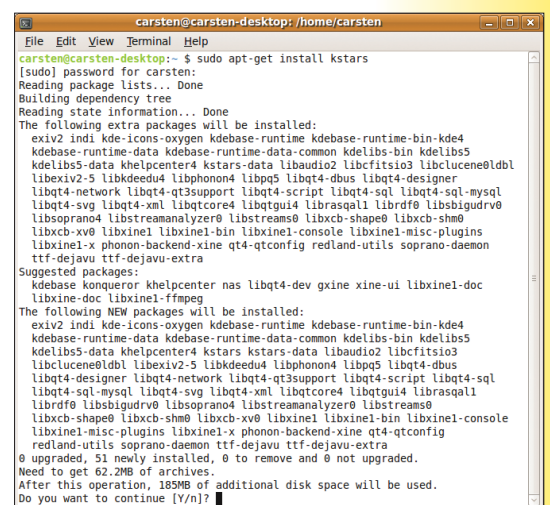


Figure 6: Installing from the command line with the *apt-get* command.

Email with Evolution

Access your email accounts, contacts, calendars, tasks, and notes with Evolution. BY MATTHEW SACKS

The Evolution email client allows a user to access email accounts, contacts, calendars, tasks, and notes all in a single location. Evolution mail is the de facto Linux email client bundled with Ubuntu 9.04. Simple to set up and configure, it can be expanded and customized for the advanced user. Evolution has been developed on one simple premise: simplicity without compromise of functionality.

Creating Email Accounts

The 2.26 release, which is shipped with Ubuntu 9.04, includes intelligent spam filtering and supports Google Contacts. To start Evolution, click on the mail icon in the top left navigation menu or go to *Applications | Internet* and click on *Evolution mail*.

For an unconfigured system, the welcome screen will appear. Clicking *Forward* advances to the *Restore from the Backup* window. If a previous Evolution archive exists that needs to be restored, check the *Restore Evolution from the backup file* box and select the archive to be restored. If no previous email accounts exist, leave this box unchecked and click *Forward*.

In this example, Gmail is the free mail service of choice. Gmail is just one of many email providers that Evolution can access to send and receive email. To access Gmail from Evolution, enable IMAP or POP access by logging

into Gmail and clicking *Settings | Forwarding and POP/IMAP | Enable IMAP* (or *Settings | Forwarding and POP/IMAP | Enable POP for all mail*). Or you can enable both for all mail. The most difficult part of configuring Evolution is getting the settings you need to enter into the client that allow you to connect to the mail server of choice. For information such as SMTP servers and POP or IMAP server addresses, check the email provider's website or call their tech support line. Typically, this information is listed on the ISP's website.

A prompt appears for identity information. First fill in the *Full Name* field with the name that will display when recipients receive messages, the *Email Address* of the account, and, optionally, the *Reply-To* address (Figure 1).

POP Email Server vs. IMAP

IMAP is chosen for the example here over POP, although most commercial and free email providers support both platforms. Evolution allows the user to access either one or both. The main difference between POP and IMAP is that, with the use of IMAP, actions executed within Evolution also are reflected on the email server. For example, if the IMAP protocol is being used and the user deletes a message within

Evolution from the IMAP account, the message also will be deleted from the email server. On a POP server, the message would only be deleted from the Evolution in box. The message would still remain on the server or be deleted from the server the first time the message is downloaded to Evolution. The difference between the use of one protocol over the other boils down to the user's preferences and what the email provider supports. The general recommendation is to use IMAP when it is available.

In the next dialog, *Receiving Email* (Figure 2), select the *Server Type* in the drop-down box. In this example, the IMAP protocol is used to connect to a Gmail account. The configuration parameters for the mail server will vary between providers. Generally, IMAP, POP, or Exchange Server (for corporate users) will be used. The configuration settings should be typed as follows for Gmail:

```
Server Type: IMAP
Server: imap.gmail.com
Username: MyUserName
```

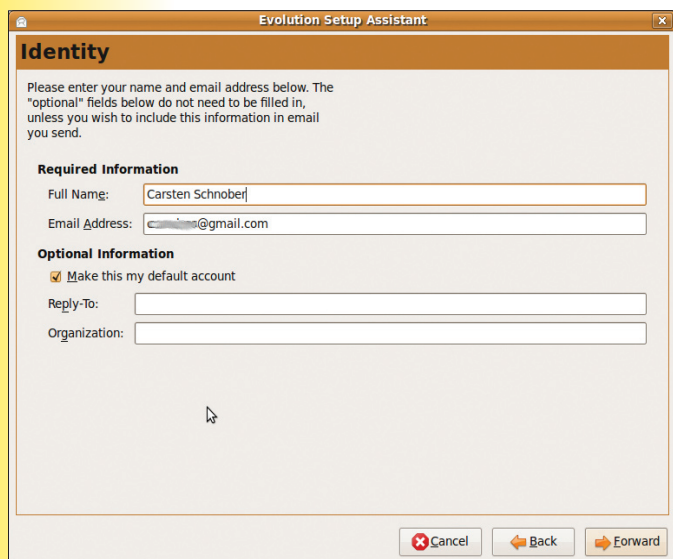
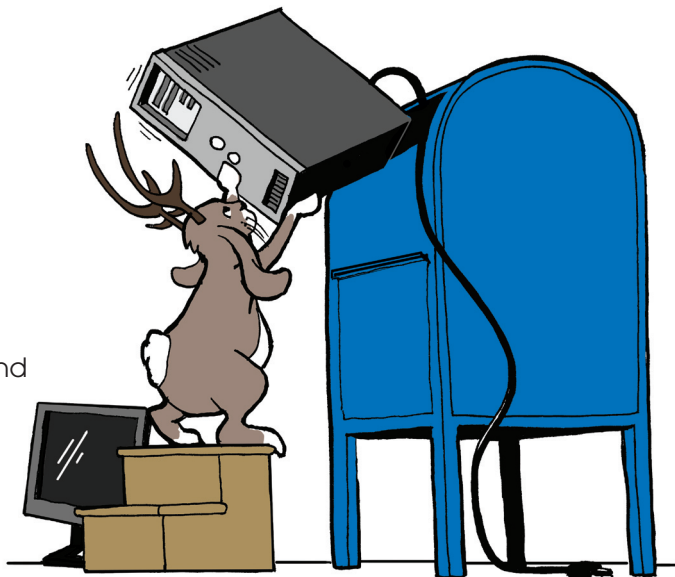


Figure 1: The *Identity* step in the Evolution Setup Assistant.

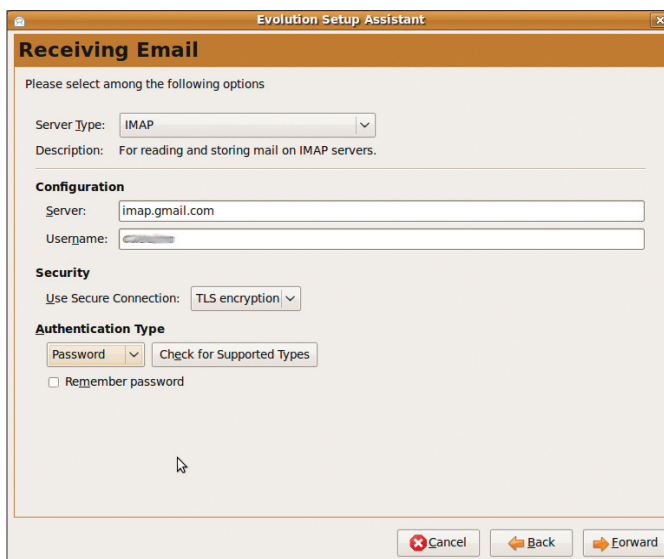


Figure 2: Setting up email as an IMAP account.



Figure 3: Specifying options for incoming mail.

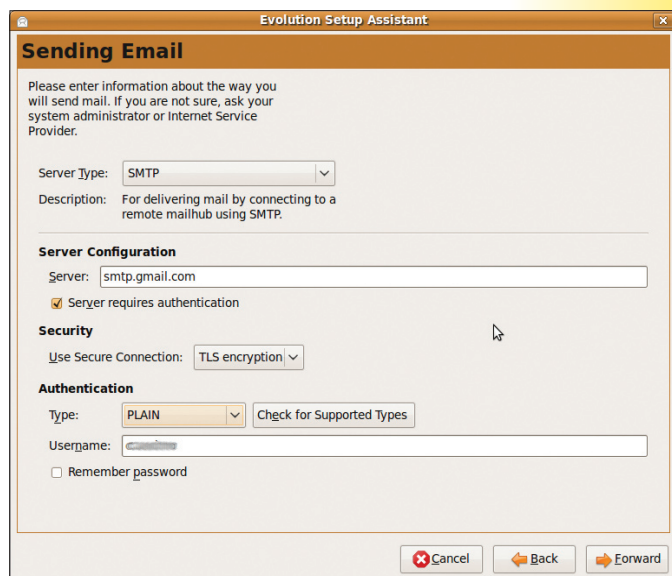


Figure 4: Specifying options for outgoing mail.

Use Secure Connction:
SSL Encryption

Next is the Receiving Options dialog box (Figure 3). Evolution checks for new messages periodically if you mark the *Check for new messages every n minutes* box. If you mark *Check for new messages in all folders*, it will download messages for all folders present, including the Sent Items mailbox and any other custom folders you have created for the email account.

- In the Options section of this dialog:
- *Apply filters to new messages in INBOX on this server* allows message filtering on the basis of sender, subject, or content of an email message. Configuring how junk mail is filtered will be covered later on.
 - *Check new messages for Junk contents* enables the SpamAssassin or Bogofilter plugin to apply spam-filtering capabilities to messages received.
 - *Automatically synchronize remote mail locally* instructs Evolution to download the contents of the email and “open” the envelope. Otherwise, IMAP by default only downloads the message headers, which is essentially the equivalent of the envelope of a physical piece of mail.

In the Sending Email dialog (Figure 4), in which you configure the outgoing mail server options, the configuration for Gmail is:

Server Type: SMTP
Server: smtp.
gmail.com
Server Requires
Authentication:
Checked
Use Secure
Connection: SSL
Authentication
Type: Login
Username:

MyUserName
Remember
Password: Checked

At this point, you can continue to the Account Management dialog box and enter the name of your account as it should appear in the Evolution email client. To select the time zone in the next window, click on the appropriate region of the world map or find the region in the *Selection* drop-down menu.

In the final Evolution Setup Assistant screen, click *Apply* to accept the changes. Now Evolution is fully configured to send and receive email with Gmail or other IMAP/POP accounts.

Most corporate mail users use Microsoft Exchange Server as their back-end mail server. Exchange accounts automatically synchronize contacts, calendars, memos, and tasks, and it typically integrates well with Linux desktops; however, with Evolution integrating the two, it is painless.

Syncing with Exchange

Evolution connects to Exchange through the Outlook Web Access (OWA) protocol, a webmail service of Microsoft Exchange Server. OWA publishes a webmail interface for the user [1], and Evolution can use this nifty interface to import the Exchange information.

To set up an Exchange account in Evolution, you need to add a new account with *Edit | Preferences*. The Evolution Preferences dialog box is where most of the configuration information for Evolution is input. First, click on the *Mail*

Accounts button on the left, then on the *Add* button on the right (Figure 5).

At this point, the Evolution Account Assistant appears. Before continuing on to the Receiving Email window, you need to supply the information for the next window (name, email address, etc.). For the *Server Type*, select *Microsoft Exchange*, then enter a *Username* and the *OWA URL*, and click *Authenticate* to make sure the settings are correct and that Evolution can connect to the Exchange server.

If no error message appears, Evolution has connected successfully to Exchange, and you can continue to the Receiving Options window. Here, you specify the interval at which Evolution checks for new messages, as well as the Global Catalog server name, which is required to retrieve contacts from an address book and access Exchange Free/Busy service when scheduling meetings. If this information is not available, you should contact your Exchange system administrator for it; otherwise, it will default to the Exchange Server information passed through OWA.

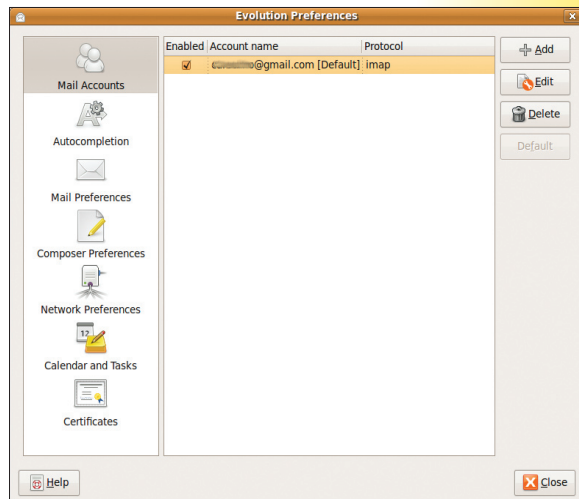


Figure 5: The Evolution Preferences dialog box.

Next is the Account Management window. Here, you enter the name of the account as it should appear within Evolution and click *Forward* and *Apply* to save the changes. Evolution begins syncing email, calendars, and contact information with Exchange. Depending on the size of the mailbox, this could take a few hours. A status message shows synchronization progress at the bottom of the Evolution Mail window.

Once the data are synchronized, any changes made within Evolution will be synced automatically to the Exchange server.

Spam Filtering

SpamAssassin is open source spam-filtering software maintained by the Apache Group. Evolution has a plugin in Ubuntu 9 that makes SpamAssassin the default spam-filtering module. It is powerful because of its predefined filtering rules, heuristic learning engine (Bayes' theorem), and integration into a worldwide network of spam identification servers.

By default, SpamAssassin is not installed with Ubuntu 9, so to use the plugin that ships with Evolution, it will need to be installed by Ubuntu's package management system. To open a command-line terminal, navigate to *Applications | Accessories | Terminal*. At the command prompt, type:

```
sudo apt-get install spamassassin
```

A prompt appears that requires root access to install software packages on Ubuntu. The dependencies are displayed, with a prompt for the user to type Y to continue:

```
msacks@linuxpromag:~$ sudo apt-get install spamassassin
[sudo] password for msacks:
msacks:
After this operation, 6947kB of additional disk space will be used.
```

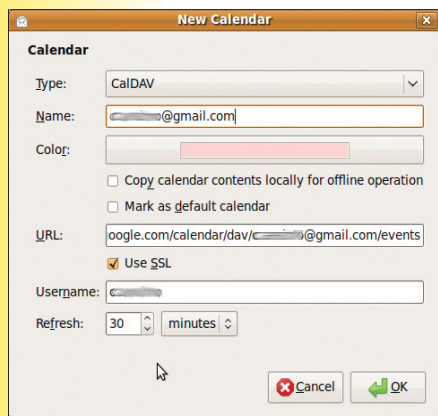


Figure 6: To synchronize a Google calendar, you should preferably use the CalDAV plugin.

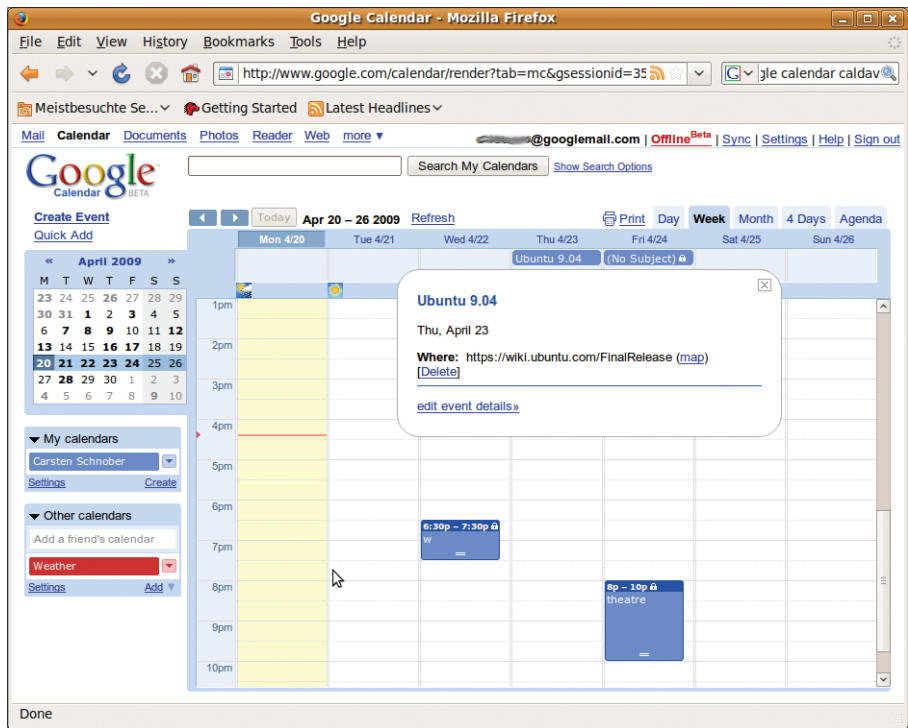


Figure 7: Google Calendar reflects entries made in the Evolution Calendar.

```
Do you want to continue [Y/n]? Y
```

At this point, the dependent software packages are downloaded and installed. Once SpamAssassin is installed, run the command

```
spamassassin -V
```

to test for successful installation. If it is installed properly, the output should indicate the SpamAssassin version and the version of Perl that was installed:

```
msacks@linuxpromag:~$ spamassassin -V
SpamAssassin version 3.2.5
running on Perl version 5.10.0
```

Now that you have enabled SpamAssassin, you must turn it on and enable it by default. To enable SpamAssassin, edit a configuration file as follows by typing

```
sudo gedit /etc/default/spamassassin
```

or

```
sudo vi /etc/default/spamassassin
```

(see Listing 1). To enable SpamAssassin, modify `ENABLED=0` to `ENABLED=1` and make sure that `CRON=1` to turn on nightly updates. To turn on SpamAssassin, open another terminal window and type:

```
msacks@linuxpromag:~$ sudo /etc/init.d/spamassassin start
Starting SpamAssassin Mail Filter Daemon: spamd.
```

To turn off SpamAssassin, enter:

```
msacks@linuxpromag:~$ sudo /etc/init.d/spamassassin stop
```

By default, SpamAssassin is enabled on restart and bootup from now on unless it is disabled by `ENABLED=0` and a computer restart.

SpamAssassin offers a wide array of tuning and configuration options on which an entire book could be written, but reading an entire book is not required to make use of its powerful and precise spam-filtering features. Simply installing SpamAssassin, turning it on, and enabling it within Evolution is all you need to get the spam-filtering power of SpamAssassin's real-time, worldwide blacklists.

Enabling SpamAssassin Within Evolution

Now that the SpamAssassin process is running in the background, you must configure Evolution to make use of its features. By default, Evolution installs the SpamAssassin plugin on Ubuntu 9.04. To verify that it is installed, go to the *Edit | Plugins* menu in Evolution and make sure the *Spam-Assassin junk plugin* is checked.

In Evolution, you configure this by clicking *Edit | Preferences*. Then click on the *Mail Preferences* button and the *Junk* tab in the top right-hand corner.

Selecting *SpamAssassin* in the *Default junk plugin* drop-down menu should show "Spam-Assassin plug-in is available and the binary is installed." To make sure the changes take effect, restart the Evolution client.

To test the effectiveness of the SpamAssassin plugin, disable it by unchecking *SpamAssassin* in *Edit | Plugins*. Then, send some spam to an account in the Evolution client. Now try turning on the plugin again and see if messages are automatically moved to the Spam folder when the plugin is enabled. SpamAssassin is very effective and accurate at filtering out spam.

The Evolution mail client supports a vast array of calendar formats, including Webcal (e.g., Apple iCal, CalDAV, Microsoft Exchange, and Google Calendars). Web calendars can be published to the web and the calendar information then accessed from Evolution. A personal calendar, or any number of public calendars, can be aggregated within Evolution. The calendar integration feature has many options for retrieving calendar information.

Certain web calendars, such as iCal and Google Calendar, are designed to be edited from within their respective services, rather than from within Evolution. Although Evolution has a Google calendar plugin, use CalDAV instead for its full compatibility with Google Calendar and increased stability in the current Evolution release (Figure 6).

Access Calendar

To access a new Google Calendar or web calendar, click *File | New | Calendar*. For this example, I will connect to a Google Calendar.

First, choose the calendar *Type*. As mentioned above, use *CalDAV* for Google, although *Google* is also present in the list. Next, choose a *Name* and a *Color* and set the *Refresh* interval. Google requires the connection to be encrypted, which is strongly recommended in any other case, too. To activate encryption,

check *Use SSL*. In the *URL* field, enter the calendar's address.

To synchronize with your Google calendar, enter *caldav://www.google.com/calendar/dav/username@gmail.com/events* into the *URL* field, replacing *username* with your Google Mail login. The same entry usually needs to be entered in the *Username* field (Figure 6). When you click *OK*, the calendar should appear in the calendar navigation pane under Google.

Now when new entries are input into Google Calendar (Figure 7) or iCal, the same appointments will appear in the Evolution Calendar (Figure 8).

Importing Contacts

Contacts can be imported from Thunderbird and Outlook in CSV format or from another Evolution contact list. Evolution also offers direct synchronization with Palm Pilot devices. Only the Microsoft Exchange feature of Evolution allows automatic synchronization of contacts. In most instances, a simple export to a file from the old mail client and an import of the file into Evolution is all it takes to capture an old contact list.

To import contacts into Evolution from Outlook, open Outlook/Outlook Express on the source computer and click *File | Import and Export*. Now select *Export to a file* and choose comma-separated values before saving the file and transporting it to the computer that is running the Evolution client.

Next, open Evolution on the destination computer and click *File | Import*, then select *Import a single file* and select the file containing the Outlook contacts in the following dialog. Finally, choose *Outlook CSV or Tab* and choose an import folder, then click *Apply*.

Conclusion

The Evolution email client can handle most of the functions that commercial email clients do,

and with better junk mail filtering built in. The features list is extensive, and Evolution's ease of use surpasses that of most commercial email clients. The Evolution email client supports a variety of mail, calendar, and contact formats, and the number of formats only grows with every release. ■

Listing 1: SpamAssassin Defaults

```
01 # /etc/default/spamassassin
02 # Duncan Findlay
03
04 # WARNING: please read README.spamd
    before using.
05 # There may be security risks.
06
07 # Change to one to enable spamd
08 ENABLED=1
09
10 # Options
11 # See man spamd for possible
    options. The -d option is
    automatically
12 added.
13
14 # SpamAssassin uses a preforking
    model, so be careful! You need to
15 # make sure --max-children is not
    set to anything higher than 5,
16 # unless you know what you're
    doing.
17
18 OPTIONS="--create-prefs
    --max-children 5 --helper-home-dir"
19
20 # Pid file
21 # Where should spamd write its PID
    to file? If you use the -u or
22 # --username option above, this
    needs to be writable by that user.
23 # Otherwise, the init script will
    not be able to shut spamd down.
24 PIDFILE="/var/run/spamd.pid"
25
26 # Set nice level of spamd
27 #NICE="--nicelevel 15"
28
29 # Cronjob
30 # Set to anything but 0 to enable
    the cron job to automatically
    update
31 # spamassassin's rules on a nightly
    basis
32 CRON=1
```

Info

[1] OWA: <http://office.microsoft.com/en-us/outlook/HA010860351033.aspx>

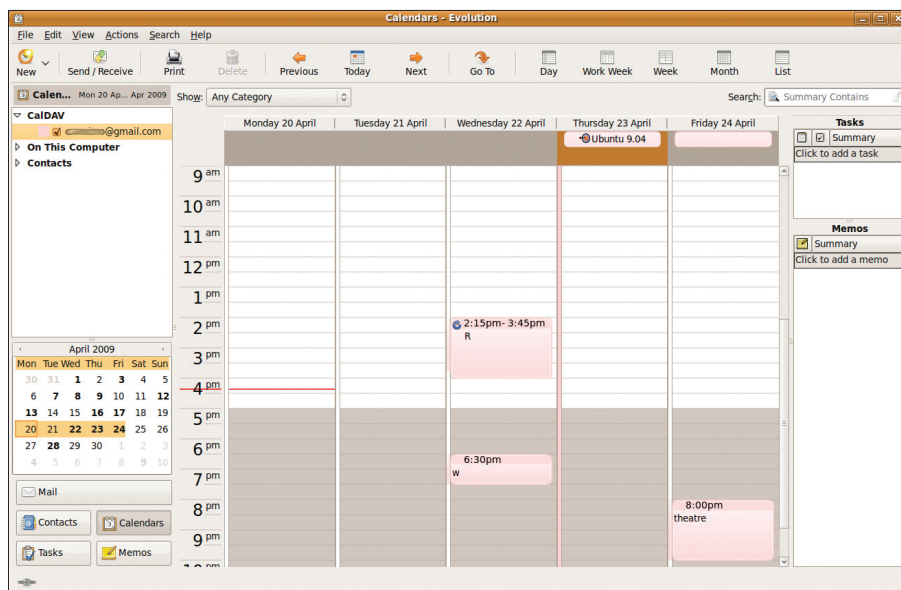


Figure 8: The Evolution Calendar reflects entries made in Google Calendar.

Taking Notes with Tomboy

Record your thoughts with the Tomboy note-taking tool. BY DMITRI POPOV

The Tomboy note-taking tool helps you create and manage personal notes. Tomboy notes are the equivalent of yellow sticky notes, except that they are stored and automatically organized on your computer. If you are running an Ubuntu flavor without Tomboy, such as Xubuntu, installing the application is easy with the Synaptic package manager or by executing the `apt-get install tomboy` command in the terminal. After you've installed Tomboy, you can launch it by choosing *Applications | Accessories | Tomboy Notes*.

To access Tomboy's commands, click on the Tomboy icon in the taskbar. The sample note that appears automatically contains a few starting points, and you can use that note or create an empty note to experiment with Tomboy's features. Each Tomboy note acts as a mini-text editor with its own toolbar that includes buttons for the most important functions (Figure 1). The *Search* button opens the Search All Notes window, where you can search your notes. To find notes containing a specific word or text fragment, enter it in the search field and press Enter for a list of matching notes.

Clicking on any note opens it in a separate window, with all occurrences of the search string highlighted for your convenience. Then you can use the *Previous* and *Find Next* buttons at the bottom of the window to jump between the found instances. To discard the search information and use the note as usual, click on the close button near the toolbar.

To perform a quick search in the current note, select the word or text fragment you want to find, right-click on it, and choose *Find in This Note | Find*. Although Tomboy doesn't support tagging notes, you can add keywords to each note and then use the *Search* feature to find notes containing a specific keyword. For example, you can add the keyword "tasks" to all your to-do lists and then use the *Search* feature to find them.

Links

Clicking the *Link* button in the toolbar allows you to link any selected word to a new note. This seemingly simple feature adds a whole new dimension to Tomboy; for example, you can keep all your appointments in one note



and then link each item on the list to a separate note. In this way, you can link the "Tomboy article" item in the to-do note to another note containing relevant links and other resources, and the "Dentist" to-do item can be linked to the note with your dentist's address.

Tomboy also adds a nifty twist to the linking feature: when you link a selected word or text fragment, the system automatically links other similar words and snippets. For example, if you link the word "Tomboy" to a note, the system auto-

matically converts all occurrences of this word into links to the note.

WikiWords

Besides manual linking, Tomboy supports so-called WikiWords, which convert any Camel-Case word into a link to a non-existing note.

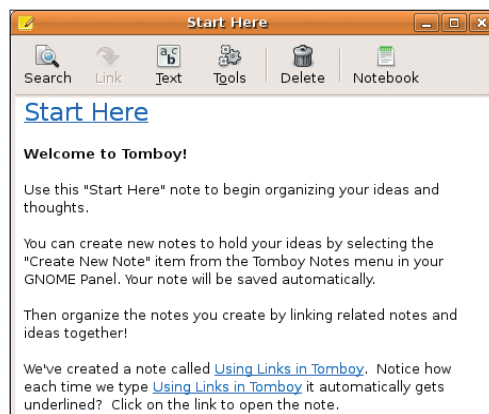


Figure 1: Each Tomboy note acts as a mini-text editor.

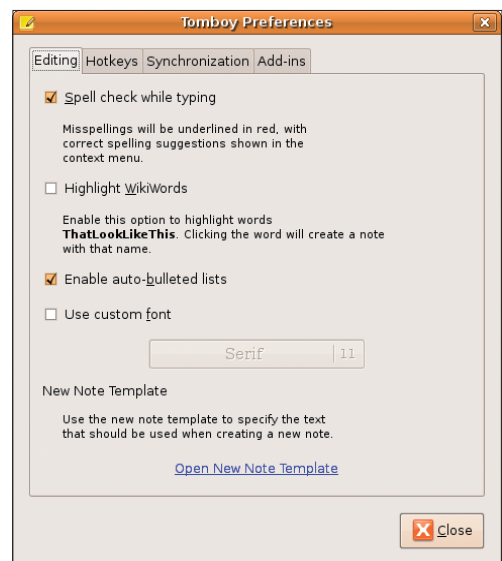


Figure 2: The Preferences window allows you to customize Tomboy's settings.

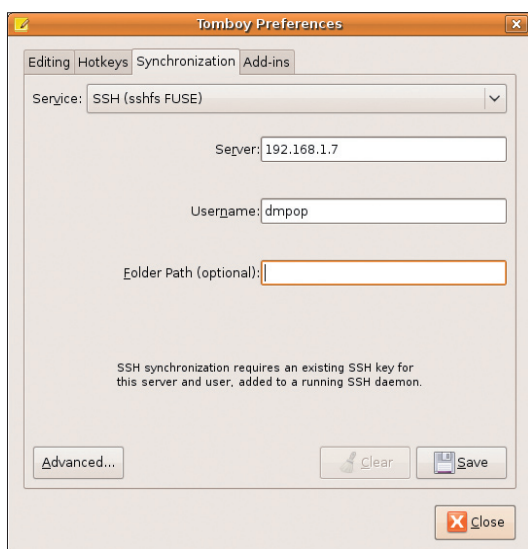


Figure 3: Tomboy can sync your notes with another machine, server, or local directory.

Because this technique is used widely in wikis, you might already know how this feature works and how to use it. By default, the Wiki-Word recognition feature is turned off, so you must first enable it in the Preferences window (Figure 2; see also “Preferences” below).

The *Text Properties* button offers a few essential text-formatting options, such as bold, italic, strikethrough, and highlight. Tomboy also supports different text sizes, as well as simple bulleted lists that come in handy when using Tomboy as a no-frills outliner or a tool for keeping track of to-do items. The *Tools* button provides access to several actions you can perform on the current note; for example, you can synchronize your notes and export them in HTML format. The *Tools* button also contains the *What links here* command, which displays a list of notes that link to the current note; this can be useful for backward navigation.

The *Delete* button doesn't require any explanation, so I'll move straight to the last button in the toolbar, which allows you to place a note in an existing notebook. To make use of this feature, you must set up at least one notebook with *Notebooks | New Notebook*. Once you've created a notebook, you can add the current note to it, or you can create a new note directly in the notebook by choosing *Notebooks | New "Notebook name" Note*. Grouping notes to notebooks makes searching easier because the search feature allows you to search through all notes or only notes in a particular notebook.

Preferences

The Preferences window, which you can access by right-clicking on the Tomboy icon, allows you to tweak Tomboy's settings. Under the

Editing tab, you can enable automatic spell check and the WikiWords feature.

If you don't like the default font, you can choose another. To do this, tick the *Use custom font* checkbox and select the font you like. Here, you can also specify the template used for every new note. The *Hotkeys* tab lets you enable and specify keyboard shortcuts for common tasks so you can create and manage notes without leaving the keyboard (Figure 5).

With the use of options available in the *Synchronization* tab, you can configure the way Tomboy synchronizes notes (Figure 3). This feature can be especially useful if you use Tomboy on several machines and want to keep your notes in sync. Moreover, you can use the synchronization feature to back up your notes. By default, Tomboy allows you to sync notes with a directory on the local hard disk or a WebDAV server. However, you also can perform synchronization with another machine via SSH, but first you must enable *SSH Sync Service Add-in* from the *Add-ins* tab (Figure 4).

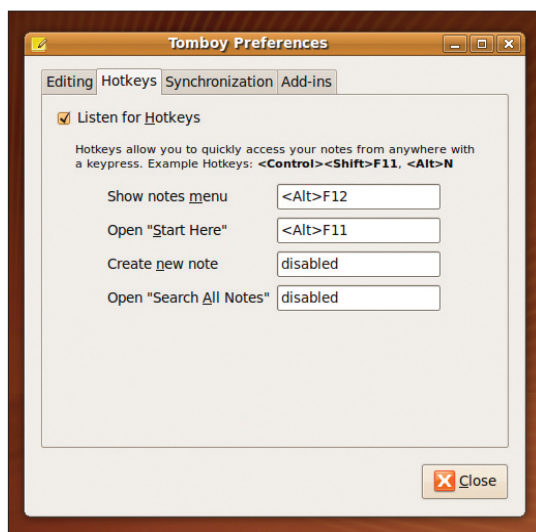


Figure 5: Create keyboard shortcuts in the *Hotkeys* tab.

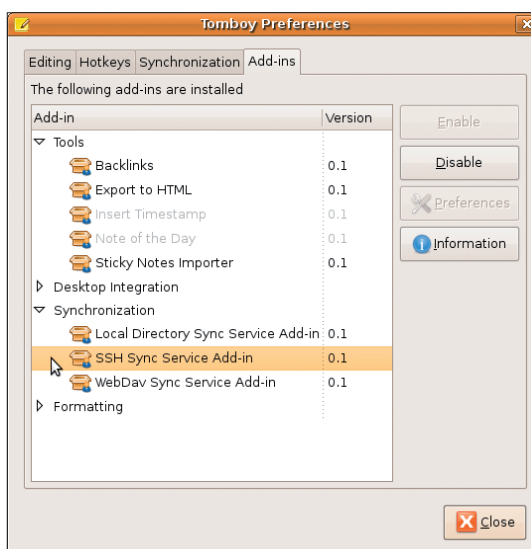


Figure 4: With the use of add-ins, you easily can extend Tomboy's functionality.

Once the add-in is enabled, select the *SSH (sshfs FUSE)* item from the *Service* drop-down list and fill out the required fields.

Tomboy Add-ins

The *Add-ins* section allows you to manage Tomboy's add-ins, which are small modules that expand the application's default functionality. Tomboy comes with a handful of officially supported add-ins, and a list of add-ins maintained by third-party developers is available on the project's wiki [1]. At least two add-ins are particularly

useful: Blogposter and Reminder.

As the name suggests, Blogposter [2] allows you to post notes to your blog, thus turning it into a simple blog publishing tool. The add-in can publish notes to Wordpress, Blogger, LiveJournal, or any blogging system that supports the APP protocol.

The Blogposter add-in is available through the Ubuntu repositories. Search for Blogposter in the Synaptic package manager to install the package. After you restart Tomboy, open the Preferences window and enable the *Post note to your blog* add-in in the *Add-ins* section. Then press the *Preferences* button and fill out the required fields.

Reminder [3] is another add-in from the same author. By adding a date and time to a note, Reminder allows you to turn the note into a reminder. For example, if you add the string *!2011/01/01 07:00* to a note, this note pops up on your screen on January 1, 2011, at 7:00am.

Reminder also understands strings such as *!Monday*, *!next monday*, and *!Monday @ 2h am*. The exclamation mark in front of each reminder entry is used to differentiate the actual reminder from regular date and time entries in the note. Downloading the *tomboy-reminder.dll* binary and moving it to the *.tomboy/add-ins* directory is the easiest way to install Reminder. Then restart Tomboy and enable the add-in in the *Add-ins* section of the Preferences window. ■

Infos

[1] Tomboy add-ins: <http://live.gnome.org/Tomboy/PluginList>

[2] Blogposter add-in: <http://flukkost.nu/blog/tomboyblogposter/>

[3] Reminder: <http://flukkost.nu/blog/tomboy-reminder/>

Open Source Office Suite

If your needs extend beyond word processing to spreadsheets, databases, and presentation software, we'll show you how to ramp up your productivity with the OpenOffice suite. BY DMITRI POPOV

The **OpenOffice suite** is a comprehensive collection of productivity tools for the home or business user. To find the word processor, spreadsheet, and presentation tools, select the Ubuntu Applications menu and choose *Office*.

Introduction to Writer

For many users, daily computing means writing and editing documents in one form or another, so a decent word processor is an essential tool for their work. Although other word processors are available for Linux, such as AbiWord and KWord, they are no match for OpenOffice.org (OOo) Writer. This heavy-duty word processor can be used for pretty much anything: From simple letters and invoices to books and mail merge documents – Writer handles these tasks equally well.

Users familiar with other word processors will feel at home with Writer. It sports a rather conventional interface (some might call it dated), and all its basic features are immediately available via the main toolbar. The Formatting toolbar provides quick access to the formatting options, such as font, font size, alignment, numbered and bulleted lists, and so on. In other words, even if you are completely new to OOo Writer, you can start using it right away. However, behind Writer's simple interface hides a rather powerful application con-

taining a slew of clever features. Writer, for example, includes a built-in PDF export feature that allows you to generate a PDF version of your document in a matter of a few clicks (Figure 1). To do this, choose *File | Export as PDF*. This opens the PDF Options dialog window that allows you to tweak the available export options. Besides PDF, OpenOffice supports a number of other formats, including Rich Text Format (RTF) and Microsoft Word 2003. The latter can be particularly useful for people who need to exchange documents with Microsoft Office users. Keep in mind, that although OOo Writer usually does a commendable job of handling the Word format, it can occasionally stumble on some complex documents containing a lot of formatting. Using OpenOffice's export feature, you can also save Writer documents as MediaWiki-formatted text files – a boon for users who use wikis for web publishing.

Like any word processor worth its salt, Writer supports version tracking, on-the-fly spell checking, footers, headers, footnotes, tables, and other essential features. Users interested in mail merge will be pleased to learn that Writer comes with a powerful yet user-friendly Mail Merge Wizard. The Frames feature can help you to lay out complex documents, and you even have a Media Player you can use to play media files and insert them into documents. All these are indispensable tools, but two features really make Writer stand out from the crowd: styles and master documents. Moreover, you can expand the capabilities of Writer and other OpenOffice applications with extensions.

Work in Writer with Style

Styles is one of the features that make OOo Writer a rather powerful word processing application. In fact, if you really want to get the most out of



Writer and take your word processing skills to the next level, you should learn how to use styles (Figure 2). The Styles feature allows you to specify paragraph and character formatting properties. For example, you can create a paragraph style called *My_bold_head* (or edit a default paragraph style) that uses the Bitstream Vera 9pt Bold font. The next time you have to format a bold head, you simply apply the *My_bold_head* to it.

But you might ask: "Why bother creating a separate style if I can just use the *Bold* button on the Formatting toolbar?" Well, say you have a 99-page document with dozens of words and text fragments in bold. Now, what if you need to change all heads to bold italic? Without styles, you have to go through the entire document and manually change every single head individually. Needless to say, this is a lot of work. With styles, performing this operation is simple: Simply open the *My_bold_head* style for editing, choose the *Font* tab, select *Bold*

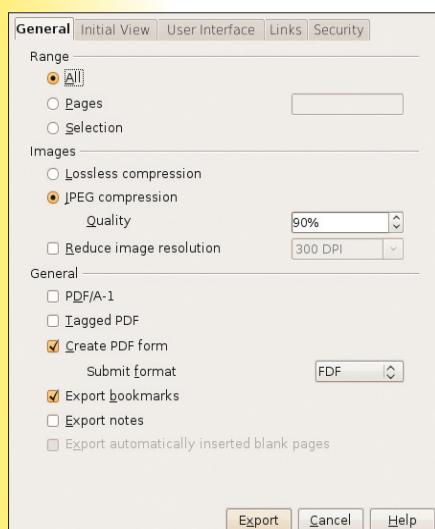


Figure 1: Save your Writer documents in PDF format with ease.

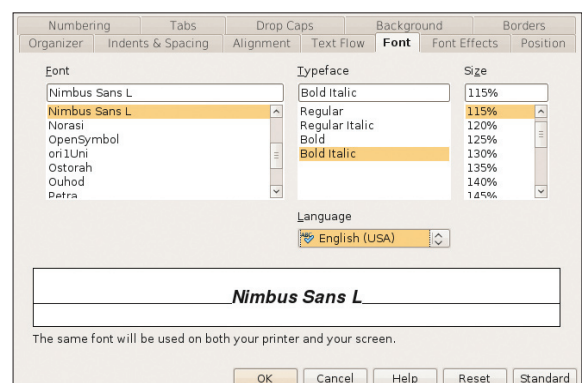


Figure 2: Tweaking styles in Writer.

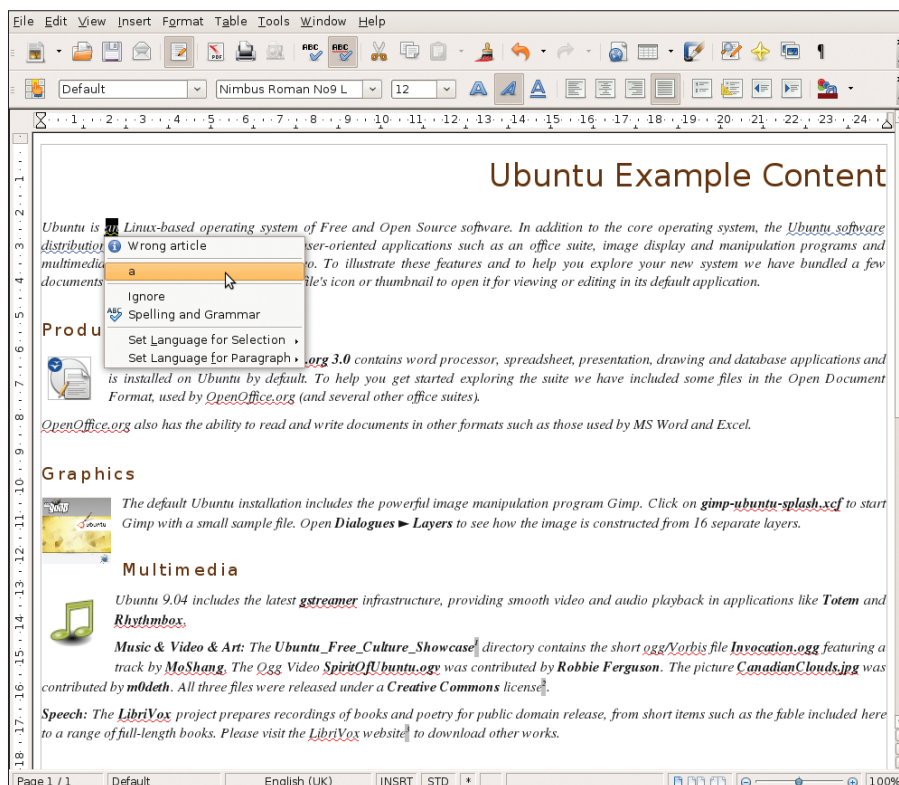


Figure 3: Checking grammar with LanguageTool.

Oblique in the *Typeface* list, press *OK*, and you are done.

The *Styles* feature offers a few other creature comforts. The *Organizer* tab in the *Style* dialog window allows you to specify which style should follow the current one (the *Next Style* option) and which style inherits its properties (the *Linked with* option). For example, you can create a special style for the document header and set the *Next Style* option to the *Text body* paragraph style. The next time you type a header in your document and press *Enter*, Writer automatically switches to the *Text body* style. Using the *Linked with* option, you can link several styles together. For example, you can link the *Heading 1* style with *Heading 2*, so the latter style inherits the properties of the former. If you later change, for example, the font in the *Heading 1* style, this change will be automatically applied to the *Heading 2* style and any style linked to it. In this way, you don't have to go through each style and change its properties manually. To keep tabs on styles, Writer sports the *Stylist* tool, which can be evoked by pressing the *F11* key or by choosing *Format | Styles and Formatting*. Styles can be a bit intimidating in the beginning, but once you've learned the ropes, they can save you a lot of time and make you look like a real OOo Writer pro.

Mastering the Master

Styles is not the only feature that makes Writer a perfect word processing tool for complex documents. When working on a long document, like a book or research paper, you can

save individual chapters as separate *.odt* files and assemble them in a master document. The master/subdocuments concept is not new. This feature has been available in Microsoft Word for a while, but it was (and probably still is) so buggy that even Microsoft recommended you stay away from it. The master documents in

Writer work like a charm, so if you are planning to write your next book in OpenOffice, you should consider using the master/subdocuments capabilities of Writer. Creating a master document in Writer couldn't be easier: Just choose *File | New | Master Document*, and Writer creates an empty master file ready for use. Then you can use the *Navigator* tool (press *F5* or choose *Edit | Navigator* to open it), which allows you to add subdocuments to the master and manage them with ease.

Extending Writer

As in Firefox, you can add new features to OpenOffice applications with extensions. The best place to find them is the official extension repository [1]. Here, you can find extensions for all OpenOffice applications, as well as rate them and comment on them. Installing OpenOffice extensions is rather straightforward. First, with the Synaptic package manager, download the desired extension, which also requires that the Ubuntu *openoffice.org-java-common* package be installed. Then launch OpenOffice, choose *Tools | Extension Manager*, and select the *My Extension* section. Now press the *Add* button and select the downloaded *.oxt* file. Once the extension is installed, restart OpenOffice to enable it.

Which extensions you choose to install on your system depends largely on your needs, but you might want to try at least two extensions: *LanguageTool* [2] and *OpenOffice.org-2GoogleDocs* [3].

Although OpenOffice comes with a competent spell checker and thesaurus, it lacks a

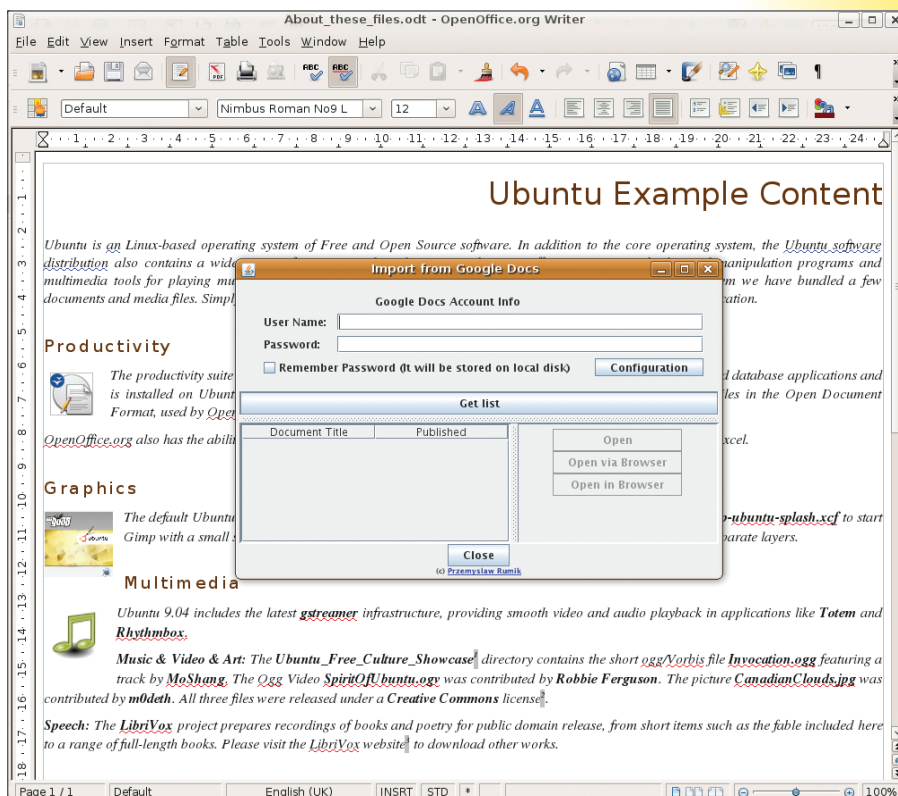


Figure 4: Use the *ooo2gd* extension to upload your documents quickly to Google Docs.

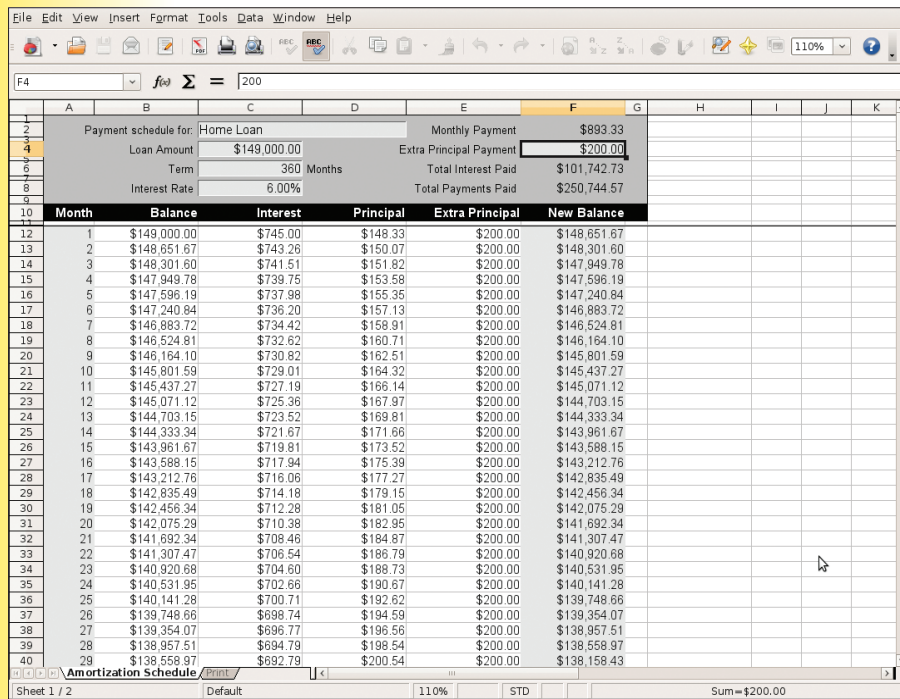


Figure 5: OpenOffice.org Calc in action.

grammar module, which many users find rather useful. If you are among them, then try LanguageTool (Figure 3). It might not be as powerful as the grammar checker bundled with Microsoft Office, but it does let you perform basic grammar checking with ease, and it supports several languages.

Right from the start, Google Docs (Docs) has been a useful complement to OpenOffice. Because this web-based office suite supports the OpenDocument format, you can send documents back and forth between OpenOffice and Google Docs. Although this method works fine when you only have to upload or download documents every now and then, if you often need to transfer dozens of documents between OpenOffice and Google Docs, it can quickly become a nuisance. The OpenOffice.org-

2GoogleDocs (ooo2gd) extension solves the problem by allowing you to upload OpenOffice documents to Google Docs, as well as open existing Docs files in Writer (Figure 4). The upload feature supports not only ODT documents, but also SXW, DOC, and RTF. Moreover, ooo2gd can send spreadsheets in ODS, XLS, and CVS formats and as PPT and PPS presentations. The download part of the ooo2gd extension is still in the experimental stage, and for now, you can only open Google Docs in Writer as plain text files.

Use of the extension couldn't be easier. When installed, it adds a submenu to the File menu containing the commands *Export to Google Docs* and *Import from Google Docs*. Alternatively, the Zoho service is also supported. Simply choose a command you need

TIP

Instead of entering the status manually each time, click on a status cell and press the Ctrl+D key combination. This gives you a drop-down menu containing all previously entered items. To make the invoice manager even more useful, you can add a feature that marks overdue invoices. To do this, simply add a Payment due column (C) and add a function to the cells in column D (see Figure 8).

and enter your username and password, and the extension does the rest.

■■■ Introduction to Calc

When it comes to number crunching and charting, the OOO Calc spreadsheet application is your best friend (Figure 5). Calc has pretty much all the features you would expect to see in a competent spreadsheet application. For starters, Calc offers a wide range of formatting options. Similar to Writer, Calc supports styles, which makes it easier to format your spreadsheet documents. Using the Ranges (*Data | Define Range*) feature, you turn selected cells into a named group. For example, you can turn a column containing sales figures into a range called "Sales" and the column containing costs into a "Costs" range. Then, you can use the ranges to perform natural language operations on them – for example, "Sales - Costs." Calc also supports conditional formatting, which applies a certain style to a cell depending on a specified condition. The spreadsheet application comes with a huge collection of functions, from something simple like *LEN* (which calculates the length of a text string) to something more esoteric like *YIELD-MAT* (which returns the annual yield of a security that pays interest at maturity). Working with functions can be a bit of a chore, but Calc comes with the Function Wizard, which makes

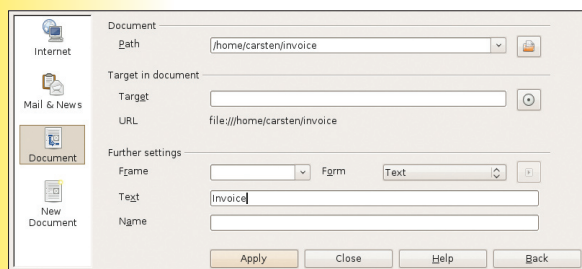


Figure 6: Insert Hyperlink dialog window.




Figure 7: Conditional Formatting dialog window.

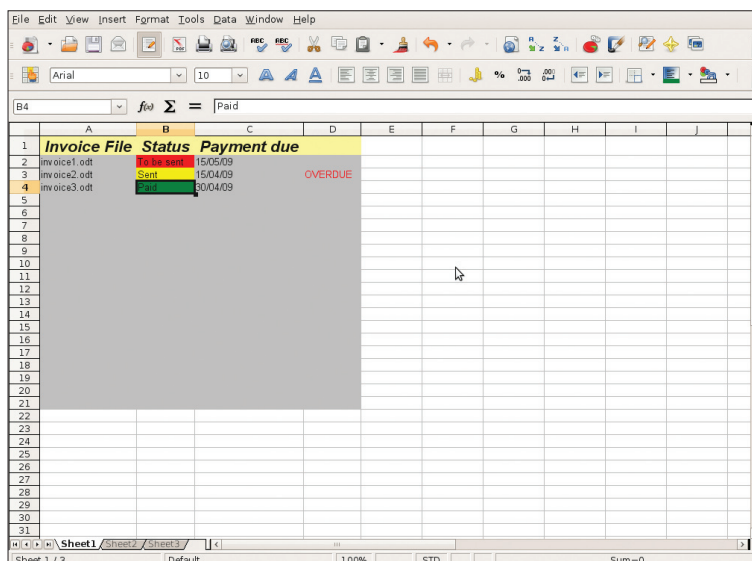


Figure 8: Invoice manager.

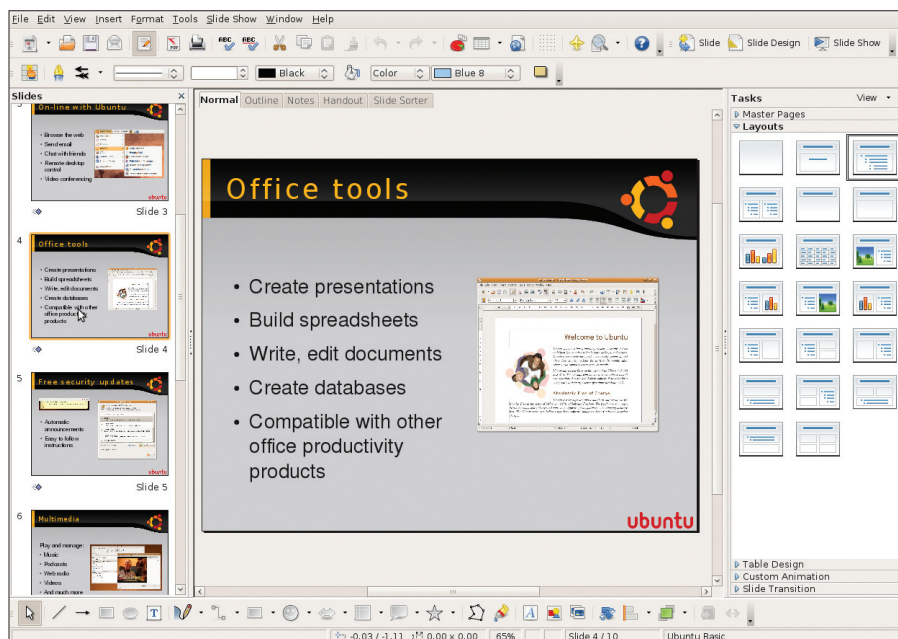


Figure 9: OpenOffice.org Impress.

it significantly easier to design simple and complex functions alike. The spreadsheet application can also pull data from external data sources (*Insert | Link to External Data*). With this feature, you can, for example, create a spreadsheet that pulls stock prices from a website and uses these data to generate charts. Speaking of charts, the latest version of Calc comes with a completely redesigned Chart module, which not only makes the process of creating charts more intuitive but also adds new chart types, such as donut and exploding donut. Calc's DataPilot (*Data | DataPilot*) feature is similar to Microsoft Excel's pivot tables, and it allows you to cross-tabulate, summarize, and convert data into meaningful information.

Creating a Simple Spreadsheet to Manage Invoices

To come to grips with Calc's basic features, I'll create a simple spreadsheet that can help manage invoices stored on a computer as Writer documents. The first order of business is, of course, to create a new spreadsheet.

To insert a link to a specific invoice, click in cell A1 and choose *Hyperlink* from the Insert menu. Select the *Document* option in the Hyperlink dialog window, then enter the path to the invoice in the *Path* field (Figure 6). In the *Text* field, enter a descriptive name for the invoice – for example, the invoice number and title. If you leave this field blank, Calc will insert the full path into the cell. To finish, click *Apply* and then *Close*.

Now that all of your invoices are in one place, how do you know which ones need to be mailed or have already been paid? To solve this problem, you need to add a status column with color codes. First, use the Stylist to create three cell styles: *To be sent*, with a red background; *Sent*, with a yellow background; and

Paid, with a green background. Also, you can choose whatever background color you like.

Now click on cell B1 and choose *Conditional Formatting* from the Format menu. In the Conditional Formatting dialog window, define the three conditions, as shown in Figure 7. The specified conditional formatting will be used to apply one of the defined cell styles depending on the cell's content.

To do the same for all, or at least some, of the cells in column B, click on the formatted cell B1, click on the black handle in the lower right corner (the cursor changes to crosshairs), hold the mouse button down, and drag the handle down.

Next, click on cell D1 and press the Function Wizard button. Use the Function Wizard

to enter the following function: `=IF(AND(C2<TODAY()-5;B2="Sent");"OVERDUE";" ")`.

This means that if the payment due date is less than the current date minus five days and the status of the current invoice is "Sent," then the invoice will be marked "OVERDUE." Use the formatting trick described above to apply the function to other cells in column D, and your invoice manager is ready to go (Figure 8).

Introduction to Impress

If you need to create a presentation that delivers your message with style, OOo Impress is here to help you. Impress' interface is divided into three panes: the Slides pane to the left, the Task pane to the right, and the working area in the middle (Figure 9). The working area also contains several tabs that you can use to switch between different sections. For example, in the *Notes* tab, you can add notes to your slides, whereas the *Slide Sorter* tab allows you to change the order of the slides. The Tasks pane consists of five sections: *Master Pages*, *Layouts*, *Table Design*, *Custom Animation*, and *Slide Transition*. The Master Pages section contains so-called master slides – special templates that you can apply to all slides in your presentation. As you might have guessed, the Slide Transition section can be used to specify transitions between slides, and Impress offers quite a few effects to choose from.

Impress offers the whole gamut of drawing tools that let you design advanced presentations with minimum fuss. Export formats for presentations include PDF and Flash, and the latter makes it easy to embed a presentation into a web page. Impress ships with a rather meager collection of presentation templates, but you can download and install the Professional Template Pack extension [4], which adds a few additional templates.

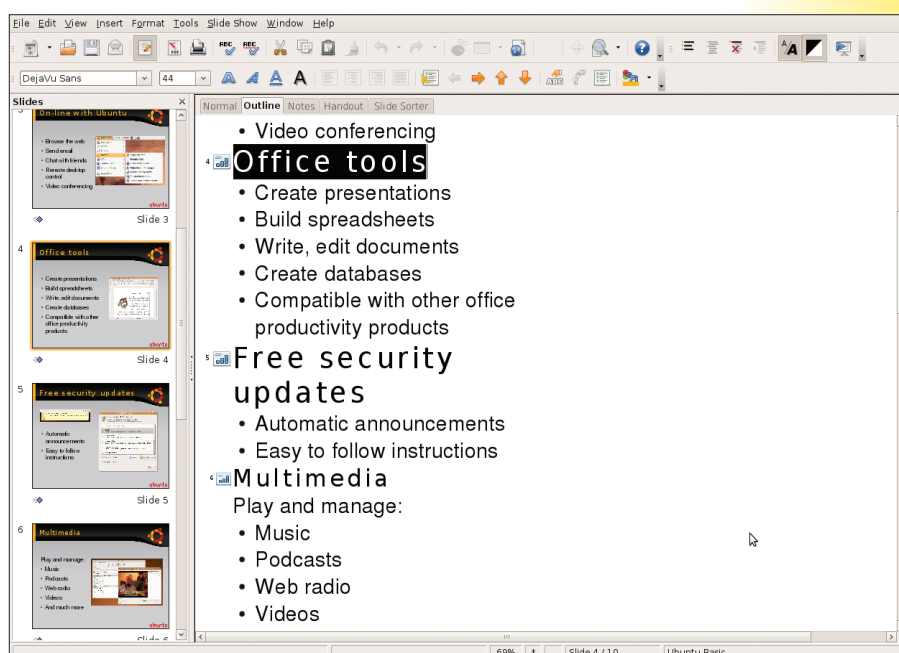


Figure 10: Working with the Outline tab.

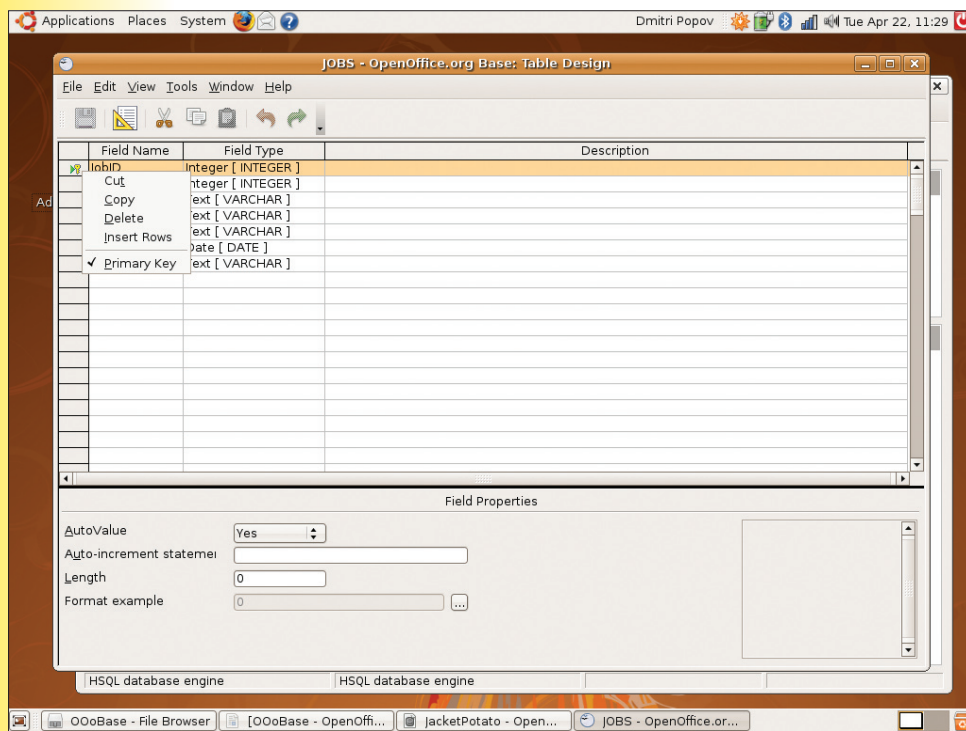


Figure 11: Setting a primary key.

Although Impress' features and interface can be a bit overwhelming in the beginning, using the application to create presentations is not that difficult. If you choose *File | New | Presentation*, OpenOffice launches the Presentation Wizard, which helps you configure the basic settings of your presentation. First, you have to choose the template and slide design you want to use for your slides. Then, you specify the output format for the final presentation and choose between overhead sheet, paper, slide, and other options. The wizard also lets you specify a slide transition that will be used throughout the entire presentation. If you want the presentation to run automatically, you can also specify the duration time for slides and pauses between them. Once you are satisfied with the settings, press *Create*, and the wizard generates a presentation.

Of course, if you prefer to create a presentation from scratch, you can skip all the described steps and simply press the *Create* button. This generates a blank presentation that you can tweak to your heart's content. If you choose to go this way, the first thing you might want to do is choose a master slide for your presentation. This will apply a specific layout to all the slides you create, thus saving you a lot of work. To choose a master slide, click on the Master Pages section and choose the master you like from the Available for Use category.

Impress offers two ways of creating slides and adding content to them. The first and most obvious way is to right-click somewhere in the Slides pane and choose *New Slide*. Then you can use the main working area to add bulleted points, graphs, and other stuff to the slide. Once the slide is ready, you create a new slide

and work on it. Alternatively, to manage your presentation as an outline, click on the *Outline* tab in the working area (Figure 10). Each time you type a header and press Enter, Impress adds a new slide. Using the promote and demote buttons in the Formatting toolbar, you can manage text as outline items. To speed things up, use the Tab key to demote the current item or Tab+Shift to promote it. The Outline mode allows you to concentrate on the content of your presentation. You can tweak the layout and overall appearance later.

Timing is everything in presentations, so Impress includes the Rehearse Timings feature. With this feature, you can adjust timings for each slide. First choose *Slide Show | Rehearse Timings*. This launches the presentation, and you should see a timer in the left corner at the bottom of the slide. Say to your imaginary audience what you want to communicate about the first slide. Then when you are ready to go to the next slide, click on the timer, and Impress saves the timing for the current slide. Now re-

peat this step for other slides in your presentation.

Introduction to Base

OpenOffice comes with its own database application called Base. Although it is a rather complex application, once you've mastered its basics, you can create pretty much anything: from a simple recipe database to sophisticated project management solutions.

Installing OpenOffice Base

By default, Ubuntu ships without OOO Base, so before you can start creating databases, you have to install the Base application on your machine. To do this, click on *Applications | Add/Remove* and select the *Office* section. Now enter *openoffice* in the search field, locate the entry *OpenOffice.org Database* in the list of found items, and tick the checkbox next to it to mark it for installation. Then press the *Apply Changes* button to install all the necessary packages.

Using OpenOffice Base

The best way to learn the ropes is to create a working database solution in OOO Base. For example, you can build a simple database solution that will help you to keep track of your jobs, how many hours you spend on each job, companies (i.e., your customers), and contact persons for each company. Launch OOO Base and choose *File | New | Database*, then use the Database Wizard to create and save a database file. When you click the *Finish* button, OOO Base opens the main window, which is the main "working zone" from which you create your database.

The first thing you need to do is create tables and populate them with fields. Choose *Tables* in the Database pane, and click on *Create Table in Design View*. This will create a new empty table, which you will use to store job data. Then create a JobID field by typing its name in the Field Name column, and set its type to Integer (*INTEGER*).

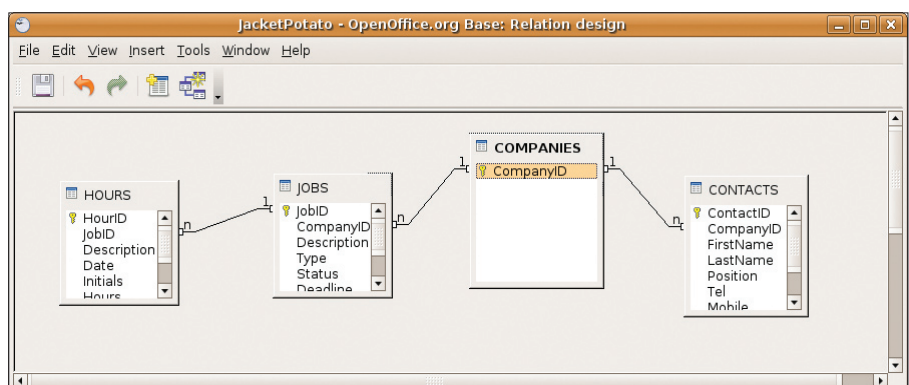


Figure 12: Defining relationships.

OOo Base offers a range of field types. The most important are numerical types (*INTEGER*, for integers – e.g., 1001, 2031, 5673; *FLOAT*, for floating-point numbers – e.g., 375.75, 12.32, 100.21); variable-length strings (*VARCHAR*, with widths specified in the Field Properties pane); dates (*DATE*); and times (*TIME*).

To set the JobID field as the primary key, right-click on the record pointer (the green triangle to the far left) and select *Primary Key* from the contextual menu. The primary key has to meet two requirements: It must not be empty, and it must have a unique value. OOo Base makes it easy to assign this property to the primary field key. Selecting *Yes* from the *AutoValue* list in the Field Properties pane (Figure 11) ensures that every new record gets a unique value (in this case, it's a number) in the JobID field.

Now add the rest of the fields and set their types. Before you save the table, make sure that the date format of the Deadline field is configured properly. Select the Deadline field, and in the Field Properties pane click on the button next to the Format example field. Select the desired format and language and click *OK*. Save the table (*File | Save* or *Ctrl+S*), and when prompted, give the table the name *JOBS*.

Now that you know how to create tables and fields, you can easily add the *HOURS*, *COMPANIES*, and *CONTACTS* tables. Just re-

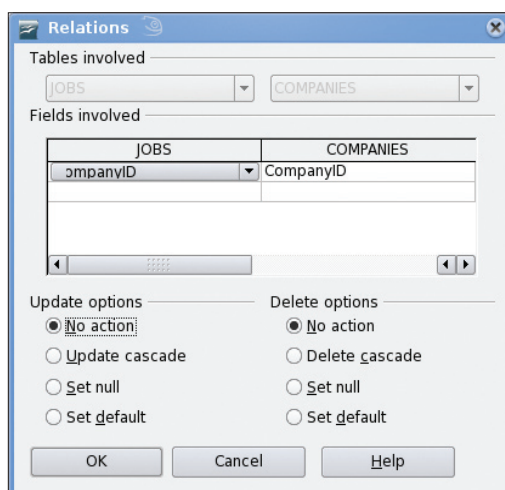


Figure 13: Setting the properties of relations.

member to set the correct format for the Date, TimeUsed, and Price fields in the *HOURS* table.

The Relation Design tools make creation of relations between tables as easy as drawing lines between primary and foreign keys (Figure 12). To define relations between the four tables, choose *Tools | Relationships* and use the Add Table dialog window to add the tables.

To establish a relation between the *COMPANIES* and *CONTACTS* tables, select the *COMPANIES.CompanyID* field, click and hold the

left mouse button, then drag a relation onto the *CONTACTS.CompanyID* field and release the button. The next step is to define the relation's properties (Figure 13). If you have a company in your database that is no longer your customer and you decide to delete it from your database, the system has to know what to do with the "orphaned" contact records. Normally, you'd want to delete them as well. Double-click on the relation between *COMPANIES* and *CONTACTS*, and in the *Delete options* column, select *Delete cascade* and click *OK*.

By the same technique, define relations between the *COMPANIES* and *JOBS* tables. Make sure, however, that the delete option is set to *No action*;

otherwise, deleting a job will automatically delete the related company. Finally, define a relation between *JOBS* and *HOURS* and set the delete option to *Delete cascade*. This will ensure that if you delete a job, you won't end up with orphaned time entries.

When the relations are defined, save them (*File | Save* or *Ctrl+S*) and close the Relation Design dialog window (*File | Close*). Now you have tables, and gradually you need to fill them with data. To view and manipulate the data, you have to create queries.

WANT TO KNOW WHAT'S UP NEXT?

**SUBSCRIBE TO LINUX
MAGAZINE PREVIEW,
OUR FREE MONTHLY
EMAIL NEWSLETTER!**

OpenOffice Base allows you to create queries with the use of either the Query Wizard or the Design View. Users that are familiar with SQL can also use the built-in SQL editor. In most cases, though, you might want to use the Design View, which allows you to construct queries in a GUI. Say you want to create a query that allows you to extract the relevant data about a job: job number, customer, job description, status, type, and deadline. Clicking on *Create Query* in the *Design View* link in the Tables section opens the Query Design window and the Add Tables dialog windows. Now you can add the JOBS and COMPANIES tables to the query with the *Add* button and click *Close*.

All you have to do now is add fields to the query, either by choosing the fields from a drop-down list in the Field cells or by dragging fields from the tables to the Field cells.

By default, the query will use the field names as labels, but you can change that to make your query more readable. For example, you can set the JobID field shown as "Job no." by typing the latter in the Alias cell (Figure 14). To save the finished query, choose *File | Save* (Ctrl+S) and give the query a name. The new query appears in the Query section, and you can launch it by double-clicking on it.

To make your database application complete, you need to build a form-based interface. Say you want to create a form that you can use to display, edit, and add jobs. Switch to the Forms section by pressing the *Forms* button, then click on *Create Form in Design View*. This will open a blank window. Click on the *Form Navigator* button on the Form Design toolbar. In the displayed Form Navigator, right-click on the *Forms* folder and choose *New | Form*. Give the form a descriptive name (e.g., *Job*), then right-click on it and select *Properties*. Under the *Data* tab, select

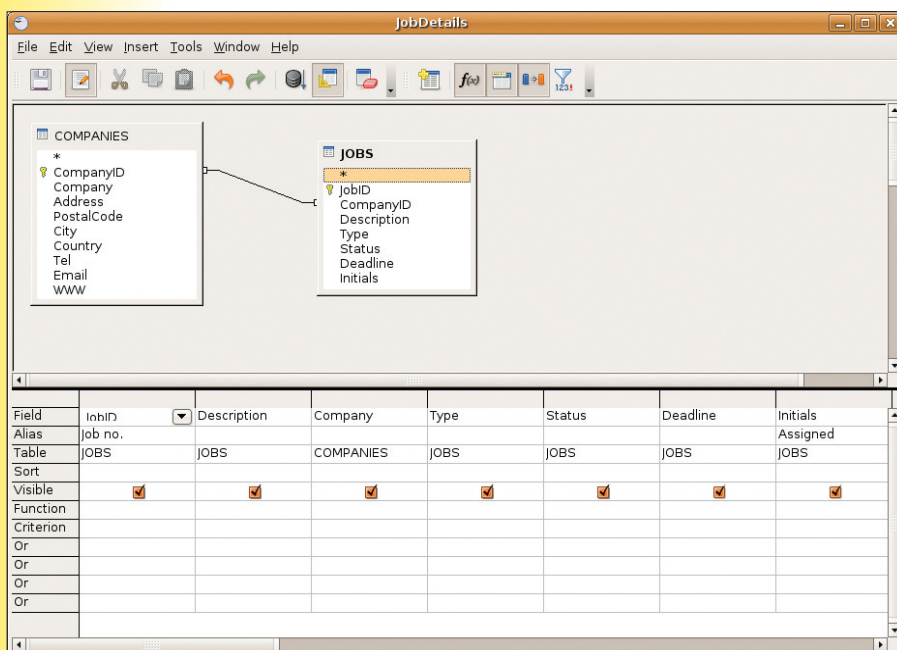


Figure 14: Creating a query.

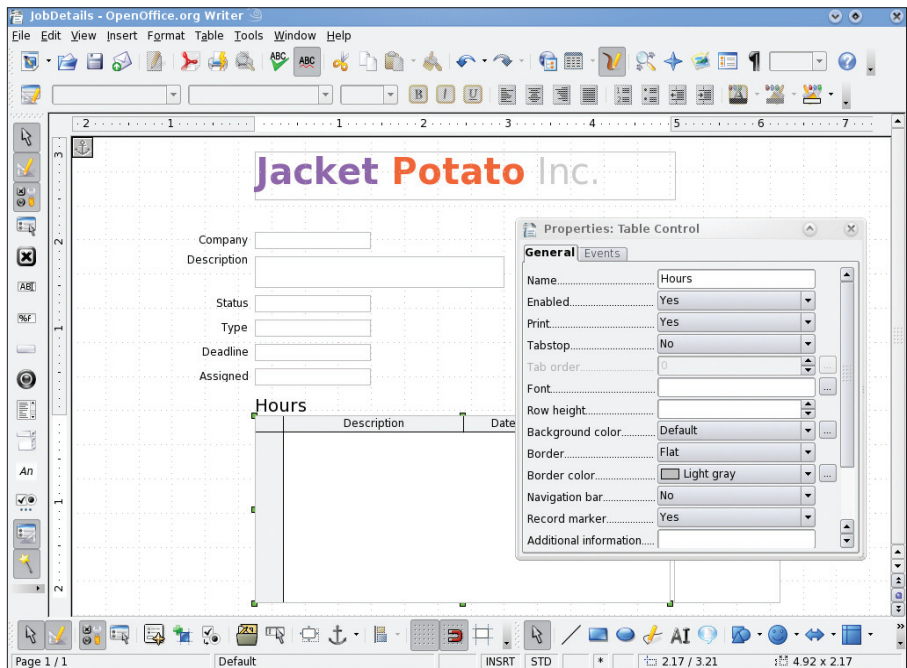


Figure 15: Designing forms in Base.

Table from the *Content Type* list, and *JOBS* from the *Content* list.

To add, for example, the *Description* text field control (which contains the job description), click on the *Text box* button on the Form Controls toolbar, and draw a text box in the form. Now double-click on the text box. Under the *Data* tab, select *Description* from the Data field list. In a similar manner you can add other fields like *Status*, *Type*, and *Deadline*. In the Properties window, you can add short help texts to the controls. To do this, click on the *General* tab, and enter the help text.

The form you created allows you to add, view, and edit data in the *JOBS* table, but you can also include data from other tables. For ex-

ample, you might want to view and register the time spent on the current job with the use of a subform.

In the Form Navigator window, click on the Job form and select *New | Form*. Give the subform a descriptive name (e.g., *Hours*), then right-click on it and select *Properties*. Under the *Data* tab, select *Table* from the *Content Type* list, and *HOURS* from the *Content* list. Then you can add form controls to the created subform. Because you can have multiple records in the *HOURS* table that are linked to a particular job, you might want to view them as a table in the Job form. For this purpose, OOO Base conveniently offers the Table Control (Figure 15). Make sure that the *Hours* subform is selected in the Form Navigator, click on the *Table Control* button and draw a table. Use the displayed Table Element Wizard to select the fields you want and press *Finish*.

When you finish creating the form, save it and close the Form Designer. Now you can double-click on the form to open and use it. To browse, add, and edit records, use the Form Navigation toolbar (*View | Toolbars | Form Navigation*). ■

Info

- [1] OpenOffice.org extensions repository: <http://extensions.services.openoffice.org>
- [2] LanguageTool: <http://www.languagetool.org>
- [3] OpenOffice.org2GoogleDocs: <http://code.google.com/p/ooo2gd>
- [4] Professional Template Pack: http://extensions.services.openoffice.org/project/SunTemplatepack_1

AUTHORS

Nick Ali	6
Jono Bacon	10
Mike Basinger	46
Florian Effenberger	56
Marcel Gagné	11
Marcel Hilzinger	67
Heike Jurzik	58
Kristian Kißling	6, 74
Andreas Kneib	53
Klaus Knopper	22
Christoph Langner	48
Achim Leitner	71
James Mohr	82
Dave Phillips	76
Dmitri Popov	50, 88, 90
Matthew Sacks	78, 84
Carsten Schnober	42, 62, 64, 67, 71
Tim Schürmann	16, 30
James Stanger	26, 36
Matthias Warkus	80

CONTACT INFO

Editor in Chief

Joe Casad, jcasad@ubuntu-user.com

Managing Editor

Rikki Kite, rkite@ubuntu-user.com

Associate Editors & Proofreaders

Carsten Schnober, Rita L. Sooby

Editorial Assistant

Heather La Bash

Contributing Editors / Test Lab Staff

Uli Bantle, Andreas Bohle, Jens-Christoph Brendel, Hans-Georg Eßer, Markus Feilner, Oliver Frommel, Marcel Hilzinger, Mathias Huber, Kristian Kißling, Jan Kleinert, Thomas Leichtenstern, Jörg Luther, Nils Magnus

Localization

Ian Travis

Layout & Cover Design

Judith Erb, jerb@ubuntu-user.com

Cartoons

K. Curtis Shontz, curt@theagilebadger.com

Cover Illustration

Pinball Werbeagentur / funkallstars, Fotolia

Advertising

www.ubuntu-user.com/Advertise

United Kingdom and Ireland

Penny Wilby, pwilby@ubuntu-user.com

Phone: +44 1787 211 100

North America

Amy Phalen, aphalen@ubuntu-user.com

Phone: +1 785 856 3434

All other countries

Hubert Wiest, anzeigen@linuxnewmedia.de

Phone: +49 89 9934 1123

Corporate Management (Vorstand)

Hermann Plank, hplank@linuxnewmedia.de

Brian Osborn, bosborn@linuxnewmedia.de

Management North America

Brian Osborn, bosborn@linuxnewmedia.com

Associate Publisher

Rikki Kite, rkite@linuxnewmedia.com

Product Management

Hans-Jörg Ehren, hjehren@ubuntu-user.com

Customer Service / Subscription**For USA and Canada:**

Email: cs@ubuntu-user.com

Phone: 1-866-247-2802

(Toll Free from the US and Canada)

Fax: 1-785-274-4305

For all other countries:

Email: subs@ubuntu-user.com

Phone: +49 89 9934 1167

Fax: +49 89 9934 1199

Ubuntu User • c/o Linux New Media •
Putzbrunner Str 71 • 81739 Munich • Germany

www.ubuntu-user.com – Worldwide

www.ubuntu-user.com.au – Australia

www.ubuntu-user.co.uk – United Kingdom

While every care has been taken in the content of the magazine, the publishers cannot be held responsible for the accuracy of the information contained within it or any consequences arising from the use of it. The use of the DVD provided with the magazine or any material provided on it is at your own risk.

Copyright and Trademarks ©2009 Linux New Media Ltd.

No material may be reproduced in any form whatsoever in whole or in part without the written permission of the publishers. It is assumed that all correspondence sent, for example, letters, e-mails, faxes, photographs, articles, drawings, are supplied for publication or license to third parties on a non-exclusive worldwide basis by Linux New Media unless otherwise stated in writing.

Printed in Germany

Distributed by COMAG Specialist, Tavistock Road,
West Drayton, Middlesex, UB7 7QE, United Kingdom

Ubuntu User ISSN 2040-8080

Ubuntu User is published by Linux New Media USA, LLC,
719 Massachusetts, Lawrence, KS, 66044, USA and Linux
New Media Ltd, Manchester, England. Company registered
in England.

Linux is a trademark of Linus Torvalds.

Ubuntu and Canonical are registered trademarks of
Canonical Ltd.

WATCH FOR US...

APPROXIMATE

UK / Europe	Aug 12
USA / Canada	Aug 26
Australia	Oct 02

ON SALE DATE

If you are looking for real answers to your Ubuntu configuration problems, or even if you just want to explore the exciting world of Ubuntu, watch your newsstand for the next issue of *Ubuntu User* magazine.

You'll learn about great free tools for the Ubuntu environment. We'll also bring you tips from the experts on scripting and troubleshooting Ubuntu systems, and we'll show how to save time and money through advanced techniques for virtualization and Windows integration.

www.ubuntu-user.com

DON'T MISS A SINGLE ISSUE!

Ubuntu User is the first print magazine created specifically for Ubuntu users. Ease into Ubuntu with the helpful Discovery Guide included in each issue, or advance your skills with in-depth technical articles, HOW-TOs, reviews, tutorials, and community reports.

SUBSCRIBE NOW!

4 issues per year for only
£ 24.90 / EUR 29.90 / US\$ 39.95

- ✓ Don't miss a single issue!
- ✓ Huge savings – Save more than 35% off the cover price!
- ✓ Free DVD – Each issue includes a Free DVD!



www.ubuntu-user.com



Now Supports **Ubuntu 9.04**
JAUNTY JACKALOPE

BACKUP LINUX SERVERS TO ANYTHING

BackupEDGE™ creates easy, verified data backups with available high strength encryption and key management. Choose the storage medium best suited to protect your critical data and systems, including...

- * **Network Attached Storage** (*intranet or Internet*)
- * Amazon S3 Internet Cloud Storage
- * Tape drives (including changers/libraries)
- * Iomega® REV® drives (including loaders)
- * Removable disk / cartridge disk storage
- * Flash drive storage
- * Standard CD and DVD writers
- * Any combination

RecoverEDGE™ bare metal disaster recovery is fully supported from any medium type. It is included with every copy of *BackupEDGE* for Linux.



What's an FTP Backup? *BackupEDGE 2.3* can stream live backups to any FTP and / or FTPS equipped storage server, NAS device or appliance, providing users with:

- * High availability and security
- * Multiple backups on-line
- * Instant access to archived files and folders
- * Bare metal disaster recovery
- * Freedom from tape management woes
- * **New!** Re-start option picks up the backup after a network link outage

Browse to: <http://www.microlite.com> right now for additional information about FTP Backups!

Only *BackupEDGE* has an advanced, standards-based storage format allowing capture and protection of ALL of your data. Features include:

- * Pathnames - up to 5,000 characters
- * Integrity checksums (file level)
- * 9 selectable compression levels (file level)
- * Archive / subset-of-archive encryption (file level)
- * Access control list (ACL) support

Transparent Media™ technology allows virtually instantaneous restore of individual files and directories from most archive resources, and the fastest possible tape positioning and restore.

Microlite Corporation has been in business since 1983 and protecting SMB and edge-of-enterprise UNIX and Linux servers since 1987. Visit our website today to download a fully functional 60-day evaluation copy of *BackupEDGE*.

<http://www.microlite.com>

