

LINUX JOURNAL

Since 1994: The Original Magazine of the Linux Community

The Ayatana Project
Improving Social Media
on the Ubuntu Desktop

Introducing Nepomuk
KDE's Semantic Desktop
Technology

The Linux Desktop

What's New with the GNOME Project and Ubuntu?



BUILD
a Portable
Custom Desktop

PLUS:
Check Out Alternate
Desktop Managers

HOW-TO:
Kexi, the KOffice Database Management Program

REVIEWED:
Giada Slim-N20,
Jolicloud and
Ubuntu 10.10
Maverick Meerkat
Amazon EC2

FEBRUARY 2011 | ISSUE 202
www.linuxjournal.com

\$5.99US \$5.99CAN 02>



0 09281 03102 4

1&1® WEB HOSTING

"Working with startup companies on a tight budget, my clients know an online presence is needed. I use 1&1 exclusively to satisfy their requests for reliable hosting at affordable rates."

Lance Ochs, www.vacantpixels.com



PROFESSIONAL WEBSITES 6 MONTHS FREE!*



*Offers valid for a limited time only. 12 month minimum contract term applies for web hosting offers. Setup fee and other terms and conditions may apply. Mobile software selection dependent on hosting package and requires a Windows PC. Domain offers valid first year only. After first year, standard pricing applies. Visit www.1and1.com for full promotional offer details. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 internet AG, all other trademarks are the property of their respective owners. © 2011 1&1 Internet, Inc. All rights reserved.



As the world's largest web host, we know the developer features you need in a hosting package!



Domains Included

All hosting packages include domains, free for the life of your package.



Unlimited Traffic

Unlimited traffic to all websites in your 1&1 hosting package.



Developer Features

Extensive language support with PHP 5/6 (beta) with Zend Framework and git version management software.



Online Marketing Tools

SEO tools to optimize your website. 1&1 Webstatistics makes it easy to monitor your progress.



Green Data Centers

We're committed to hosting your site with a minimal impact on the environment.

1&1® HOSTING PACKAGES
**6 MONTHS
FREE!***

1&1® BUSINESS PACKAGE:

- 3 Included Domains
- Private Domain Registration
- 250 GB Web Space
- UNLIMITED Traffic
- **NEW:** Version Management Software (git)
- 2,500 E-mail Accounts
- 50 MySQL Database (100 MB)
- 25 FTP Accounts
- E-mail Marketing Tool
- 24/7 Toll-free Customer Support

~~\$9.99~~
per month*

Need more domains?

.info domain only \$0.99 first year*

.com domain only \$4.99 first year*

More great offers available on our website!



Get started today, call 1-877-GO-1AND1

www.1and1.com

CONTENTS FEBRUARY 2011

Issue 202



FEATURES

DESKTOP

44

Coming Soon to Linux Desktops

Why the shift away from standard GNOME, and will Unity work for you?

Charles Olsen

48

The Second-String Desktop

Your bulky default desktop might not really be as powerful as you once thought.

Shawn Powers

54

Organize Your Life with Nepomuk

Nepomuk: hyped, mysterious and...useful?

Stuart Jarvis

60

Create Your Own Linux Desktop and Take It with You

Sometimes it would be nice to have Linux on your keyring.

Rick Rogers

ON THE COVER

- The Ayatana Project: Improving Social Media on the Ubuntu Desktop, p. 74
- Introducing Nepomuk, KDE's Semantic Desktop Technology, p. 54
- What's New with the GNOME Project and Ubuntu?, p. 44
- Build a Portable Custom Desktop, p. 60
- Plus: Check Out Alternate Desktop Managers, p. 48
- Reviewed: Giada Slim-N20, p. 36; Jolicloud, p. 39; Ubuntu 10.10 Maverick Meerkat Amazon EC2, p. 41
- How-To: Kexi, the KOffice Database Management Program, p. 68

This month's cover is a screenshot of reader Vinay Jhedu's desktop. Vinay is the winner of our recent desktop screenshot contest on LinuxJournal.com.

SPEND LESS. STORE MORE.

Aberdeen's AberNAS Network Attached Storage appliances, featuring the high performance Intel® Xeon® processor 5600 series, deliver maximum efficiency while providing the most storage capacity available today.



DENSITY

Storage from 2TB to 132TB
in a single appliance

RELIABILITY

RAID 6, Automatic RAID rebuild,
Snapshot Data Recovery and
Industry Leading 5-Year Warranty

SCALABILITY

Expand to beyond 1 Petabyte
via XDAS and JBOD units

Who gives you the best bang for the buck?

ABERNAS 160 SERIES



1U NAS up to 8TB

- Quad-Core Intel Xeon Processor
 - 3GB 1333MHz DDR3 Memory
- Available with Windows or Linux OS
- 2TB Starting at..... **\$2,795**
8TB Starting at..... **\$3,795**

ABERNAS 260 SERIES



2U NAS up to 16TB

- Quad-Core Intel Xeon Processor
 - 3GB 1333MHz DDR3 Memory
- Available with Windows or Linux OS
- 4TB Starting at..... **\$4,295**
16TB Starting at..... **\$6,295**

ABERNAS 290 SERIES



2U NAS up to 24TB

- Quad-Core Intel Xeon Processor
 - 3GB 1333MHz DDR3 Memory
- Available with Windows or Linux OS
- 12TB Starting at..... **\$5,995**
24TB Starting at..... **\$7,995**

ABERNAS 360 SERIES



3U NAS up to 32TB

- 2x Quad-Core Intel Xeon CPUs
 - 6GB 1333MHz DDR3 Memory
- Available with Windows or Linux OS
- 16TB Starting at..... **\$7,995**
32TB Starting at..... **\$10,995**

ABERNAS 560 SERIES



5U NAS up to 48TB

- 2x Quad-Core Intel Xeon CPUs
 - 6GB 1333MHz DDR3 Memory
- Available with Windows or Linux OS
- 24TB Starting at..... **\$10,495**
48TB Starting at..... **\$14,495**

ABERNAS 860 SERIES



8U NAS up to 100TB

- 2x Quad-Core Intel Xeon CPUs
 - 6GB 1333MHz DDR3 Memory
- Available with Windows or Linux OS
- 50TB Starting at..... **\$18,495**
100TB Starting at..... **\$26,995**

ABERNAS 8960 SERIES



8U NAS up to 132TB

- 2x Quad-Core Intel Xeon CPUs
 - 6GB 1333MHz DDR3 Memory
- Available with Windows or Linux OS
- 128TB Starting at..... **\$33,995**
132TB Starting at..... **\$34,995**



CONTENTS

FEBRUARY 2011

Issue 202

COLUMNS

- 20** Reuven M. Lerner's At the Forge
Communication in HTML5
- 26** Dave Taylor's Work the Shell
Dealing with Spaces in Filenames
- 28** Kyle Rankin's Hack and /
Status Messages in Screen
- 78** Kyle Rankin and Bill Childers'
Point/Counterpoint
Tablets
- 80** Doc Searls' EOF
Dysintegration

REVIEWS


- 36** Giada Slim-N20
Shawn Powers
- 39** Jolicloud
Mike Diehl
- 41** Ubuntu 10.10 Maverick Meerkat in
Amazon EC2
Bill Childers

INDEPTH

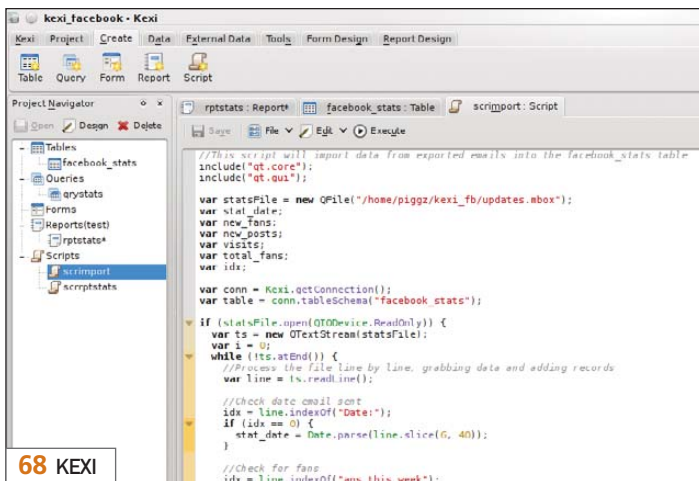
- 68** Kexi in Use
Database applications just got easier.
Adam Pigg
- 74** Social from the Start
The social-networking aspects in Ubuntu.
Jono Bacon

IN EVERY ISSUE

- 8** Current_Issue.tar.gz
- 10** Letters
- 14** UPFRONT
- 30** New Products
- 32** New Projects
- 65** Advertisers Index



36 GIADA SLIM-N20



USPS *LINUX JOURNAL* (ISSN 1075-3583) (USPS 12854) is published monthly by Belltown Media, Inc., 2121 Sage Road, Ste. 310, Houston, TX 77056 USA. Periodicals postage paid at Houston, Texas and at additional mailing offices. Cover price is \$5.99 US. Subscription rate is \$29.50/year in the United States, \$39.50 in Canada and Mexico, \$69.50 elsewhere. POSTMASTER: Please send address changes to *Linux Journal*, PO Box 16476, North Hollywood, CA 91615. Subscriptions start with the next issue. Canada Post: Publications Mail Agreement #41549519. Canada Returns to be sent to Pitney Bowes, P.O. Box 25542, London, ON N6C 6B2

More TFLOPS, Fewer WATTS

Microway delivers the fastest and greenest floating point throughput in history

Enhanced GPU Computing with Tesla Fermi

- ▶ 480 Core NVIDIA® Tesla™ Fermi GPUs deliver 1.2 TFLOP single precision & 600 GFLOP double precision performance!
- ▶ New Tesla C2050 adds 3GB ECC protected memory
- ▶ New Tesla C2070 adds 6GB ECC protected memory
- ▶ Tesla Pre-Configured Clusters with S2070 4 GPU servers
- ▶ WhisperStation - PSC with up to 4 Fermi GPUs
- ▶ OctoPuter™ with up to 8 Fermi GPUs and 144GB memory

New Processors

- ▶ 12 Core AMD Opterons with quad channel DDR3 memory
- ▶ 8 Core Intel Xeons with quad channel DDR3 memory
- ▶ Superior bandwidth with faster, wider CPU memory busses
- ▶ Increased efficiency for memory-bound floating point algorithms

Configure your next Cluster today!

www.microway.com/quickquote
508-746-7341



2.5 TFLOPS

10 TFLOPS

5 TFLOPS



4.5 TFLOPS

FasTree 864 GB/sec
Bi-sectional Bandwidth

FasTree™ QDR InfiniBand Switches and HCAs

- ▶ 36 Port, 40 Gb/s, Low Cost Fabrics
- ▶ Compact, Scalable, Modular Architecture
- ▶ Ideal for Building Expandable Clusters and Fabrics
- ▶ MPI Link-Checker™ and InfiniScope™ Network Diagnostics

Achieve the Optimal Fabric Design for your Specific MPI Application with ProSim™ Fabric Simulator

Now you can observe the real time communication coherency of your algorithms. Use this information to evaluate whether your codes have the potential to suffer from congestion. Feeding observed data into our IB fabric queuing-theory simulator lets you examine latency and bi-sectional bandwidth tradeoffs in fabric topologies.



GSA Schedule
Contract Number:
GS-35F-0431N

Microway
Technology you can count on™

LINUX JOURNAL™

Since 1994: The Original Magazine of the Linux Community

**DIGITAL EDITION
NOW AVAILABLE!**

Read it first

Get the latest issue before it
hits the newsstand

Keyword searchable

Find a topic or name
in seconds

Paperless archives


Download to your computer for
convenient offline reading

Same great magazine

Read each issue in
high-quality PDF

Try a Sample Issue!

www.linuxjournal.com/DLISSUE



LINUX JOURNAL

Executive Editor Jill Franklin
jill@linuxjournal.com

Senior Editor Doc Searls
doc@linuxjournal.com

Associate Editor Shawn Powers
shawn@linuxjournal.com

Art Director Garrick Antikajian
garrick@linuxjournal.com

Products Editor James Gray
newproducts@linuxjournal.com

Editor Emeritus Don Marti
dmarti@linuxjournal.com

Technical Editor Michael Baxter
mab@cruzio.com

Senior Columnist Reuven Lerner
reuven@lerner.co.il

Security Editor Mick Bauer
mick@visi.com

Hack Editor Kyle Rankin
lj@greenfly.net

Virtual Editor Bill Childers
bill.childers@linuxjournal.com

Contributing Editors

Ibrahim Haddad • Robert Love • Zack Brown • Dave Phillips • Marco Fioretti • Ludovic Marcotte
Paul Barry • Paul McKenney • Dave Taylor • Dirk Elmendorf • Justin Ryan

Proofreader Geri Gale

Publisher Carlie Fairchild
publisher@linuxjournal.com

General Manager Rebecca Cassity
rebecca@linuxjournal.com

Senior Sales Manager Joseph Krack
joseph@linuxjournal.com

Associate Publisher Mark Irgang
mark@linuxjournal.com

Webmistress Katherine Druckman
webmistress@linuxjournal.com

Accountant Candy Beauchamp
acct@linuxjournal.com

Linux Journal is published by, and is a registered trade name of, Belltown Media, Inc.
PO Box 980985, Houston, TX 77098 USA

Editorial Advisory Panel

Brad Abram Baillio • Nick Baronian • Hari Boukis • Steve Case
Kalyana Krishna Chadalavada • Brian Conner • Caleb S. Cullen • Keir Davis
Michael Eager • Nick Faltsy • Dennis Franklin Frey • Alicia Gibb
Victor Gregorio • Philip Jacob • Jay Kruiuzenga • David A. Lane
Steve Marquez • Dave McAllister • Carson McDonald • Craig Oda
Jeffrey D. Parent • Charnell Pugsley • Thomas Quinlan • Mike Roberts
Kristin Shoemaker • Chris D. Stark • Patrick Swartz • James Walker

Advertising

E-MAIL: ads@linuxjournal.com
URL: www.linuxjournal.com/advertising
PHONE: +1 713-344-1956 ext. 2

Subscriptions

E-MAIL: subs@linuxjournal.com
URL: www.linuxjournal.com/subscribe
PHONE: +1 818-487-2089
FAX: +1 818-487-4550
TOLL-FREE: 1-888-66-LINUX
MAIL: PO Box 16476, North Hollywood, CA 91615-9911 USA
Please allow 4-6 weeks for processing address changes and orders
PRINTED IN USA

LINUX is a registered trademark of Linus Torvalds.




AUSTIN, TX

MARCH 11-15, 2011

SXSW[®]

SOUTH BY SOUTHWEST[™]



INTERACTIVE

The South By Southwest Interactive Festival

Register to Attend

Go to sxsw.com/attend for the best available rate.

Hundreds of panels announced!

Visit sxsw.com/interactive/talks for details.

sxsw.com





SHAWN POWERS

Not Even the Droid X is Big Enough for My Desk

Some of you may argue that all desktops eventually will be on tiny screens in your pockets. To be honest, however, I like my 26" screen, and my pockets aren't nearly big enough to hold it. So, although the embedded market certainly is important, this month we're focusing on the desktop—or laptop, if that's how you roll.

Charles Olsen gives a glimpse at GNOME 3, which isn't quite ready yet, but it offers some significant changes you might want to read about. Following his article, Shawn Powers (hey, that's me!) gives an overview of the other desktop environments—possibly some you've never heard of, but many you might want to try. Sometimes you want a faster desktop, and sometimes you just want to be different from the crowd. Both are reasons to check out the available options.

Having a desktop is only the start, of course. The real work is what you do while using that desktop. Stuart Jarvis helps keep things organized with Nepomuk. Most operating systems index files for quicker searching, but Nepomuk goes one step further and integrates the process into KDE. Although I still don't suggest saving everything to your desktop folder, if that's how you do things, Nepomuk might keep you organized despite yourself. If organization is something you prefer to do manually, that's fine as well. One frustration with that scenario, however, is when you move from computer to computer. Rather than organize each computer you use, why not just take your desktop with you? Rick Rogers shows how to create your own portable Linux desktop you can take with you wherever you go.

This month, we've included plenty of useful productivity information as well. Adam Pigg shows how to use the KOffice database tool, Kexi. Databases may not be the most exciting thing you'll deal with in a day, but they often pay the bills. Rather than using a separate proprietary database format, Kexi relies on SQLite for its underlying database program. That means

any program able to deal with SQLite can deal with Kexi databases. It's great to see a program support standards, and Adam explains the ins and outs of using Kexi.

Bill Childers has his head in the clouds again this month and shows off the free 55-minute trial of Amazon's EC2 available for Ubuntu 10.10 users. Granted, 55 minutes isn't very long, but it's enough to give you a taste of the cloud with no commitment. If you discover you like EC2, but don't want to pay the monthly bill, be sure to read Bill's article. He has a solution for that as well. What he doesn't have a solution for, however, is the argument he and Kyle Rankin have about tablet computing. Whether you think tablets are the next logical step to a *Star Trek* world, or just laptops that really could use keyboards, Bill and Kyle both have good points. Perhaps I'll have to get a tablet myself—you know, for research.

We also have our regular lineup of columns, like Reuven Lerner's HTML5 discussion and Dave Taylor's help in dealing with spaces in filenames. Kyle Rankin shows how to get pop-up notifications in your terminal windows, and John Knight explores cool new programs that are fresh from the labs.

Additionally, we have a review by yours truly of the Giada N20, a sleek little ION2-based Nettop, and Mike Diehl shows off Jolicloud. We've also got info on Opendgear's new cellular router and a book on making LEGO guns. It's truly a fun issue. So, whether you think desktop computing is on its way out, or if you have a wall of 30" monitors you refuse to give up, this issue has something for you. As for me, I'm going to keep my phone in my pocket and my 26" monitor on my desk. Even with my SCOTTEVEST jacket, toting around a giant monitor doesn't sound like fun. I'll leave that chore for Bill and his tablet computer. ■

Shawn Powers is the Associate Editor for *Linux Journal*. He's also the Gadget Guy for LinuxJournal.com, and he has an interesting collection of vintage Garfield coffee mugs. Don't let his silly hairdo fool you, he's a pretty ordinary guy and can be reached via e-mail at shawn@linuxjournal.com. Or, swing by the #linuxjournal IRC channel on Freenode.net.

★ S O U T H W E S T ★ DRUPAL SUMMIT

HOUSTON, TX

JANUARY 27-28



2 DAYS
OF DRUPAL
DEVELOPMENT
BUSINESS &
DESIGN

REGISTER NOW!

WWW.SWDRUPALSUMMIT.COM

MAGNOLIA HOTEL ★ 1100 TEXAS AVE ★ HOUSTON TX 77002

See what the **Drupal Content Management System and Framework** can do for you. The event brings Drupal experts, novices, and business leaders together to share successes, explore opportunities, and learn more about why and how Drupal is making headlines across the world as a superior enterprise-level web application platform.

**REGISTER
EARLY!
VISIT US
ONLINE**



WHEN IS IT
JANUARY 27-28, 2011

WHERE TO GO
MAGNOLIA HOTEL - HOUSTON, TEXAS

HOW TO BE A SPONSOR
CONTACT JHAYS@NEOSPIRE.NET OR
KATHERINE@LINUXJOURNAL.COM

KEYNOTE SPEAKER
ANGELA BYRON, LULLABOT.COM

Angela Byron is the Drupal 7 core maintainer, recipient of the 2008 Google- O'Reilly Open Source Award for Best Contributor, and an Open Source evangelist who lives and breathes Drupal. We are so excited to welcome her to Houston!



SPONSORED BY

NEOSPIRE **LINUX**
MANAGED HOSTING JOURNAL

Lullabot



SPECIAL HOTEL GROUP RATE: \$139/NIGHT
CALL 1.888.915.1110 AND ASK FOR DRUPAL CONFERENCE GROUP
EARLY REGISTRATION IS AVAILABLE



Say a Prayer for Open Source—Literally!

I just recently finished attending a Northern Minister's Convention in Sault Ste Marie where many of the ministers were overwhelmed with costs of not just hardware, but also software that they felt necessary to purchase to maintain their congregations' technology. Upon doing some research, one particular denomination did not qualify for one of the world's largest software manufacturer's "charitable organization discount". Upon hearing this, I piped up, "What about OpenOffice.org?" It was as if I were speaking Latin! Nobody had heard of or considered it.

I have discovered over the years many pieces of software that have saved churches thousands of dollars in software expenses, such as OpenOffice.org, OpenSong (lyrics/projection/chord charts), OpenLP (lyrics and projection), SQL-Ledger (book-keeping) and many more. I believe that the charitable donation sector (aka churches) is a prime candidate to take the ultimate leap to open source and migrate to Linux as its OS of choice. When you consider the thousands of dollars a church could spend on software in a year, I would sooner see that money going to someone in a community that needs it rather than a large, monopoly-like company. So if you have a chance, pray for open source...and for the church! I will.

--
Dean Anderson

I'm a deacon at a church, and sadly, even I haven't ever heard of OpenSong or OpenLP. I'm both ashamed and now really excited to learn of them! You're right, of course. I think churches are a prime place for open-source software to fit right in. Thanks for the software advice!—Shawn

File Renaming

I enjoyed Dave Taylor's article in the November 2010 issue "Scripting Common File Rename Operations". I did find an error in the sample, where cut is used to get the extension of a file:

```
f=foo.bar.txt
echo "$f" | cut -d. -f2
```

The above will result in "bar.txt" and not "txt".

This works, however:

```
echo ${f##*.}
```

Second, I have had good experiences using co-processes in file rename operations. For example, I often need to convert filenames to lowercase:

```
paste <(ls -1) <(ls -1 | tr 'A-Z' 'a-z') | \
{ while read old new; do mv -v $old $new; done; }
```

Treating the lists of names as text opens up a lot of possibilities in file operations at reduced script complexity.

The following snippet using script generation works if there are spaces in the filenames without having to fiddle with IFS:

```
paste <(ls -1) <(ls -1 | tr 'A-Z' 'a-z') | \
gawk 'BEGIN { FS="\t"; } { printf("mv -v \"%s\" \
\"%s\"\n", $1, $2); }'
```

The output can be piped into bash to run the script. Generating the output before actually running the script makes it possible to see and catch errors before actually doing anything.

One also can use co-processes to output and increment a counter generating

a column of numbers to use in incrementally numbering files with a common base name:

```
paste <(ls -1) <(let i=0; for f in *; do echo
newFileBase_${(i++)}.${f##*}; done)
```

I look forward to future editions of Dave's article!

--
Jeffrey Brendecke

Dave Taylor replies: I appreciate the feedback on my latest column, Jeffrey! I can't say I have much experience with co-processes though. Sometimes I feel like I'm too old-school Linux/UNIX, especially since I've been using it since 1980—jeez, that's a long time!

Mutt

Regarding the article on Mutt and especially the Point/Counterpoint discussion between Kyle and Bill on this (LJ, October 2010), it's a pity neither of them seems to be aware of Muttator (vimperator.org/muttator). I would've appreciated each gentleman's standpoint on that one.

--
Ludo

D-Bus

I can't thank Koen Vervloesem enough for his wonderful article on D-Bus (LJ, November 2010). As a Linux "old-timer", it's been quite a learning experience for me to deal with newer methodologies of hardware detection and management. I ran into the world of D-Bus, Hal and udev about three years ago with my installation of openSUSE 11.0, and access to devices was well, not quite up to par. Thus, I began my immersion into this new world! Mostly, I concentrated on learning about Hal and udev, not thinking at all about D-Bus (what the heck is that, I thought), and trust me, that was way too much information for me at the time. However, I was able to put some PolicyKit rules in place successfully to accomplish my goal. I took Koen's advice and tracked down a graphical D-Bus app—what I could find for my current openSUSE 11.3 implementation

was something called KDBus. I think this is lacking *some* of the features of qdbusviewer, but wow, how cool at any rate. Now that KDE 4.x (and GNOME) also uses D-Bus for session events, any time you do anything—accessing hardware or applications in your desktop environment—you can see and explore the new D-Bus events that are fired. It's very useful and informative to track issues. Since my Linux setup still uses D-Bus, Hal and udev, seeing what events get fired and are associated with which of these hardware management layers is really very useful, and in some cases, critical to “fixing” issues that arise. Again, much thanks to Koen for this great article!

--
Kay Schenk

Keep Trying Linux

I just renewed my annual subscription and decided to add the digital edition as well. I have to say, I am exceptionally pleased with the fine work you all do at putting out *Linux Journal*, and you can bet I will continue to be a customer.

I have read often about the problems people have run into using Linux to replace Windows systems, and my advice is just to keep trying Linux and open source. I have been able to replace every Windows program with a Linux program, (KMyMoney, Evolution, OpenOffice.org, and so on) with one exception—video editing. Adobe Premiere keeps me tied to a partition I have set up to run WinXP and Premeire Elements. Maybe one day Adobe will realize it can keep a lot of customers happy and port Premiere to Linux.

My distro of choice is Ubuntu. I run Mythbuntu and Xubuntu with no problems. Thanks again for the good work, and I look forward to a new year with *Linux Journal*.

--
linuxdog977

As a longtime Final Cut Pro user, I feel your video-editing pain. Kris Occhipinti (some of his videos are on our Web site) convinced me to try Kdenlive. Although I'm still not as comfortable with it as I am with Final Cut, it does seem to be a viable alternative. It's just taking some learning on my part.—Ed.

Don't Build Your Own NAS with RAID-5

Regarding Shawn Powers' “Drobo Shmobo” piece in the November 2010 Upfront section, I hope *nobody* builds their own NAS with RAID-5. RAID-5 has a bad failure mode when rebuilding an array due to a bad disk. When a bad disk is found, you have a nontrivial chance of finding that you actually had *two* bad disks during the rebuild, causing the array to be lost. RAID-6, on the other hand, is able to rebuild an array with two failed disks.

In any event, regardless of RAID-5 or RAID-6, you should make sure you run disk checks periodically. For md-raid, for example, for md0:

```
# echo check > /sys/block/md0/md/sync_action
```

You may want to put that in root's crontab or in a script in, say, /etc/cron.weekly.

--
David N. Lombard

Although I appreciate your worries about multiple disk failures and RAID-5, for me (I wrote the article), another drive just isn't possible. I already have my workstation-size computer maxed out with drives, and if I were to sacrifice one of the drives for RAID-6, I wouldn't have enough storage. Perhaps my huge array filled with 2TB drives isn't the norm, so I'll agree if it's feasible, RAID-6 certainly makes sense.

Oh, and your disk check is right on. It's very good practice, and something Bill Childers reminded me to set up the last time I had a hard drive failure in my server!—Ed.

A Superior Technology

Puppy Studio began as an experiment. Was it possible to have a full-featured multimedia creation suite similar to Ubuntu Studio running on Puppy Linux? Puppy was fast and ran entirely in RAM, so the real-time recording capabilities could be amazing. But, would they?

My original hypothesis was correct. Puppy gave the studio vastly improved latency, by virtue of running in RAM and its small size, unmatched by any other OS. Then, I compiled a custom real-time kernel, and performance went through the roof! 3.1 is

out: www.murga-linux.com/puppy/viewtopic.php?p=458075#458075.

--
10wt3ch

Simple Success

I'm a physics teacher with a new toy. I've got a Tablet PC issued by my school. I GNU (ha!) the instant I got the thing what I was going to do with it: run Linux. We have nice IT guys who generously tolerate such behavior. They were frankly intrigued when I told them I created a Linux partition on my hard drive. One of my IT buddies said, “Most teachers aren't going to use even 20% of the capacity of these things. Let me know how it works out.” Here's what a simple Linux partition can do: play a DVD without taking ten minutes to detect the disc or freezing. Even better, there is absolutely no issue connecting to LCD projectors or detecting monitors. It is plug and play. Xarnal lets me write on PDF'd homework sets and save them to the class Dropbox file, so students can see what I did later from their phones, laptops and desktops. Need some extra homework problems? Not an issue. I can hop on the command line and use `wget -i` to download an entire MIT physics class (with PDF problem sets). Why re-invent the wheel?

Is Linux perfect? No. I don't like Linux screencasting software. I'm frankly resentful that rotating screens on my tablet PC is such a pain. I wish there were a driver for my wireless LCD connector. But, it is really hard to beat free-as-in-beer software that runs fast.

I'm really excited to continue this dual-booting experiment in the classroom. With work and time, it would be great if my entire class were run on Linux. I guess I could moan and groan that it's not totally feasible just yet, but I'm choosing to focus on the answers not the problems. It's hard not to feel grateful to *Linux Journal* and all the people who are way smarter than me who have shown me all these tools that I can use without having to think too much. Thanks guys!

--
Riaz Abdulla

It's great to hear the tablet mostly does what you need in Linux. Honestly, the cooler part for me is to hear your IT

[LETTERS]

department was okay with letting you fiddle with Linux at all! I think that attitude will help get Linux into education more than all the free software alone ever could do.—Ed.

Transparent Firewalls with Linux Series

I've enjoyed the recent series of articles by Mick Bauer in *Linux Journal* on transparent firewalls, but unfortunately, Mr Bauer has some misunderstandings about how the hardware (Linksys WRT54GL) actually works, which leads to some erroneous conclusions.

The Linksys WRT54GL is based on a Broadcom BCM5352 SoC (System-on-Chip). The key word here is system, since part of the system includes a six-port switch. Five of these ports you can see, and one is internal to the chip.

In part IV of the series (December 2010), Mr Bauer can't understand why iptables ignores inter-VLAN traffic. The reason is simple. The kernel never sees it because it is handled by the switch, and packets never make it to the CPU. Also, port 5 (the common port) is not a virtual port but real in every sense of the word. It exists in silicon and is part of the switch that connects to the CPU.

Unfortunately, full specifications for this chip seem to be available only under a non-disclosure agreement, but generic overview pamphlets are available that show this. This is also discussed on the OpenWrt forum.

--
Gus Wirth

Mick Bauer replies: *Thanks much for the clarifications and for confirming my suspicion that entire categories of the WRT54GL's packets are never seen by the kernel or CPU. Note that in at least one case, I'm guilty of oversimplification rather than ignorance: when I call port 5 "virtual", I mean only that you can't plug an Ethernet cable in to it, not that it doesn't exist in some other tangible way.*

It would be great if the OpenWrt team would address transparent

firewalling in its official documentation, or at least in a FAQ or wiki. In my lengthy research efforts for this project, key details concerning the inner workings and interactions of WRT54G hardware, OpenWrt and relevant parts of the Linux kernel were either missing from OpenWrt documentation, mentioned only in passing in on-line forums or simply nowhere to be found.

Best Linux Vendor?

I'm just reading the Readers' Choice Awards for 2010 (see the December 2010 issue), and I notice that the best Linux vendor for both laptops and desktops is Dell. I thought things must have changed at Dell since I last looked at its offerings, so I checked out what it had to offer. With great difficulty, and after a great deal of searching, I finally was able to locate some laptops that came with Linux preloaded—with Ubuntu 9.10. So, I'm confused by how Dell got a Readers' Choice Award for being the "Best Vendor".

--
lauzon

I guess you'd have to ask all the readers why they voted for Dell! In all seriousness though, my guess is the votes were based on price and compatibility, not pre-installed options. There certainly are more Linux-friendly vendors, but all too often when we order hardware anymore, we don't even bother to look at what is pre-installed. I don't think that's a good thing, but I think it happens a lot.—Ed.

PiTiVi

Just wanted to drop you guys a note and say thanks to Jono Bacon (and to LJ) for the great article on PiTiVi in the December 2010 issue. I do a video review show about comic books and graphic novels (www.youtube.com/user/The2dogknight), and I use PiTiVi to put the show together, but I never knew how to do fades. Now I do! I plan to use fades in my next episode. Love the magazine, and please keep up the great work!

--
Kevin Starkey

LINUX JOURNAL

At Your Service

MAGAZINE

PRINT SUBSCRIPTIONS: Renewing your subscription, changing your address, paying your invoice, viewing your account details or other subscription inquiries can instantly be done on-line, www.linuxjournal.com/subs. Alternatively, within the U.S. and Canada, you may call us toll-free 1-888-66-LINUX (54689), or internationally +1-818-487-2089. E-mail us at subs@linuxjournal.com or reach us via postal mail, Linux Journal, PO Box 16476, North Hollywood, CA 91615-9911 USA. Please remember to include your complete name and address when contacting us.

DIGITAL SUBSCRIPTIONS: Digital subscriptions of *Linux Journal* are now available and delivered as PDFs anywhere in the world for one low cost. Visit www.linuxjournal.com/digital for more information or use the contact information above for any digital magazine customer service inquiries.

LETTERS TO THE EDITOR: We welcome your letters and encourage you to submit them at www.linuxjournal.com/contact or mail them to Linux Journal, PO Box 980985, Houston, TX 77098 USA. Letters may be edited for space and clarity.


WRITING FOR US: We always are looking for contributed articles, tutorials and real-world stories for the magazine. An author's guide, a list of topics and due dates can be found on-line, www.linuxjournal.com/author.

ADVERTISING: *Linux Journal* is a great resource for readers and advertisers alike. Request a media kit, view our current editorial calendar and advertising due dates, or learn more about other advertising and marketing opportunities by visiting us on-line, www.linuxjournal.com/advertising. Contact us directly for further information, ads@linuxjournal.com or +1 713-344-1956 ext. 2.

ON-LINE

WEB SITE: Read exclusive on-line-only content on *Linux Journal's* Web site, www.linuxjournal.com. Also, select articles from the print magazine are available on-line. Magazine subscribers, digital or print, receive full access to issue archives; please contact Customer Service for further information, subs@linuxjournal.com.

FREE e-NEWSLETTERS: Each week, *Linux Journal* editors will tell you what's hot in the world of Linux. Receive late-breaking news, technical tips and tricks, and links to in-depth stories featured on www.linuxjournal.com. Subscribe for free today, www.linuxjournal.com/enewsletters.



Lullabot™

Learn Drupal & jQuery

FROM THE COMFORT OF
YOUR LIVING ROOM



The Lullabot Learning Series includes everything you need to become a Drupal & jQuery expert from the comfort of your living room! The videos are available in both DVD format and high-definition video download.

Purchase the videos at <http://store.lullabot.com>

diff -u

WHAT'S NEW IN KERNEL DEVELOPMENT

There's a new **dd** in town. Specifically, it's **ddpt**, written by **Douglas Gilbert**, **Mark Knibbs** and others. The **pt** is short for pass through and refers to the **SCSI pass through interface** that **ddpt** specializes in. Douglas announced the project and provided the URL: sg.danny.cz/sg/ddpt.html. The original **dd** tool, on which **ddpt** is based, is a low-level, super-powered data copier, designed (among other things) to convert data between formats recognized on different architectures.

For those of you who've been missing **Jon Masters' kernel podcast**, it's back! The feed is at podcasts.jonmasters.org/kernel/kernel.xml, and the transcripts are at www.kernelpodcast.org. It's a tough project, due to the psychotically high volume of mailing-list traffic he needs to cover. He took a break a while back as an alternative to burn-out, and he's now announced that he's back and going at it again. Have fun this time around, Jon!

Linux is famous for its high uptimes. But, nowadays people want to maximize resource usage by cramming lots of **KVM** instances onto a given hardware system, and they want each virtual machine to have either a high uptime or at least a quick recovery time. **Fernando Luis Vazquez Cao** and other folks are trying to improve KVM's high-availability powers by triggering a recovery event whenever a KVM instance goes down. It's much more straightforward to do this for a virtual machine than for an actual hardware computer. When a physical computer goes down, it also might bring down whatever triggering mechanism is supposed to start the recovery process. When a virtual machine goes down, there's much

less danger that the enclosing process might be damaged, so the trigger always can be initiated.

There's a new **IRQ Subsystem maintainer!** Or at least, it's the same old guy, but now it's in the **MAINTAINERS** file. **Joe Perches** recently posted a patch listing **Thomas Gleixner** as the **IRQ maintainer**. This was based on the fact that Thomas submits almost 50% of **IRQ** patches, while the next highest contributor submitted only 13%. **IRQs** (interrupt requests) are very quick responses to events on the system. They let you keep browsing the Web while behind the scenes everything keeps going on the way it's supposed to. And, now there's a maintainer for all that.

Miklos Szeredi announced a new **union filesystem** he'd been working on. Union filesystems are neat, because if you have two different filesystems with a similar directory structure, you can use a union FS to overlay them both on top of each other so they appear to be a single directory tree. File **A** in **/home** on one filesystem and file **B** in **/home** on the other filesystem, for example, then would appear to be files **A** and **B** in **/home** on a single unified filesystem. There also are various interesting implementation problems, such as what to do if both filesystems have a file **A** in **/home**. Which one does the user see? Miklos' union filesystem would assign one filesystem to be the "upper" one, and the other to be the "lower" one. If any files were duplicated across both filesystems, only the one in the "upper" filesystem would be visible to the user. **Neil Brown** discussed this and more in the documentation he contributed to the project, in response to Miklos' announcement.

—ZACK BROWN

They Said It

It has become appallingly obvious that our technology has exceeded our humanity.

—Albert Einstein

Technology is like a fish. The longer it stays on the shelf, the less desirable it becomes.

—Andrew Heller

The most overlooked advantage of owning computers is that if they foul up there's no law against whacking them around a bit.

—Eric Porterfield

My biggest fear regarding cloud computing is that it sounds disturbingly close to "SkyNet".

—Shawn Powers

The future masters of technology will have to be light-hearted and intelligent. The machine easily masters the grim and the dumb.

—Marshall McLuhan

LinuxJournal.com

As you flip through the awesome articles in this month's *Linux Journal*, you probably will think a lot about your own personal desktop preferences, your favorite desktop applications, or perhaps some tips and tricks you like to use to make your life a little easier. While you are pondering these things, I'd love for you to consider writing about them for LinuxJournal.com and sharing them with our on-line audience. We always are looking for tips and tricks, how-tos and other articles for our on-line publication, and if you'd like to share your knowledge with our Web readers, please send an e-mail with your idea to webeditor@linuxjournal.com. I'd love to hear your ideas and see your byline on LinuxJournal.com!

Don't forget, there is a lot more great information about desktop Linux at LinuxJournal.com. Visit www.linuxjournal.com/tag/desktop to find more articles to help you get the most out of your Linux machine. See you there!

—KATHERINE DRUCKMAN

NON-LINUX FOSS

If you've ever taken notes in Linux, chances are you've at least tried Tomboy. It's installed by default on most systems, and it's a rather impressive application for jotting down quick notes. One of its newer features is the ability to sync notes between computers. One of the benefits of that feature is that the note-syncing works cross-platform! Because there are native versions of Tomboy for both Windows and OS X, syncing between Linux and non-Linux computers should be a breeze.

The syncing feature is relatively new to Tomboy, so it's possible you'll run into a few snafus, but for cross-platform, open-source note-syncing, it's worth the effort. Download installers for all three platforms at projects.gnome.org/tomboy, although if you're running Linux, it's probably already installed.

—SHAWN POWERS



Netflix and Linux, It's a DRM Shame



By the time this goes to press, the Boxee Box by D-Link will be in the wild. I've written about this little guy before, but one of the exciting and frustrating last-minute announcements is that Netflix would be one of the content partners. When you add the Boxee Box to the Roku devices, and even a few televisions, you have a whole lineup of Linux-powered devices streaming Netflix!

Surely that is reason to celebrate. Unfortunately, it's a bittersweet celebration, because although many Linux-powered devices can indeed stream Netflix content, Linux desktop users still can't stream to their computers. Why, you ask?

Simple: DRM.

In order to approve Netflix streaming in embedded Linux devices, Netflix has included some sort of proprietary video decoder that obviously works under Linux. Why Netflix won't release that proprietary, DRM-laden software to the greater Linux community is a mystery. Is Netflix afraid it will get reverse-engineered so people can watch its streams for free? Well, unfortunately for Netflix, that probably will happen someday. If Netflix would release a proprietary, binary-only player for Linux users, the motivation to reverse-engineer would be close to zero.

So Netflix, do you want to protect your evil DRM? My suggestion is to release your proprietary binary-only software quickly. I warn you, nothing hacks like a spurned geek. Oh, and fellow geeks? Why not drop Netflix PR a note regarding your frustration. The PR page doesn't specifically have a section for DRM complaints, but it seems to me it's a *public relations* problem: www.netflix.com/ContactPR.

—SHAWN POWERS

TS-WIFIBOX-2 A Complete Solution for 802.11g WiFi Applications



qty 1 \$185



Powered by a
250 MHz ARM9 CPU

- Low power (3.2 watts), fanless
- Power via 5-12 VDC, USB, PoE (opt.)
- 64MB DDR-RAM
- 256MB ultra-reliable XNAND drive
- Micro-SD Card slot
- RS-232, RS-485, CAN, RTC
- Header with SPI and 11 DIO
- 480Mbit/s USB, Ethernet, PoE option
- Boots Linux 2.6.24 in < 3 seconds
- Un-brickable, boots from SD or flash
- Customizable FPGA - 5K LUT
- Optional DIN mountable enclosure

Ideal for gateway or firewall, protocol converter, web server, WiFi audio, and unattended remote applications

- Over 25 years in business
- Never discontinued a product
- Engineers on Tech Support
- Open Source Vision
- Custom configurations and designs w/ excellent pricing and turn-around time
- Most products ship next day

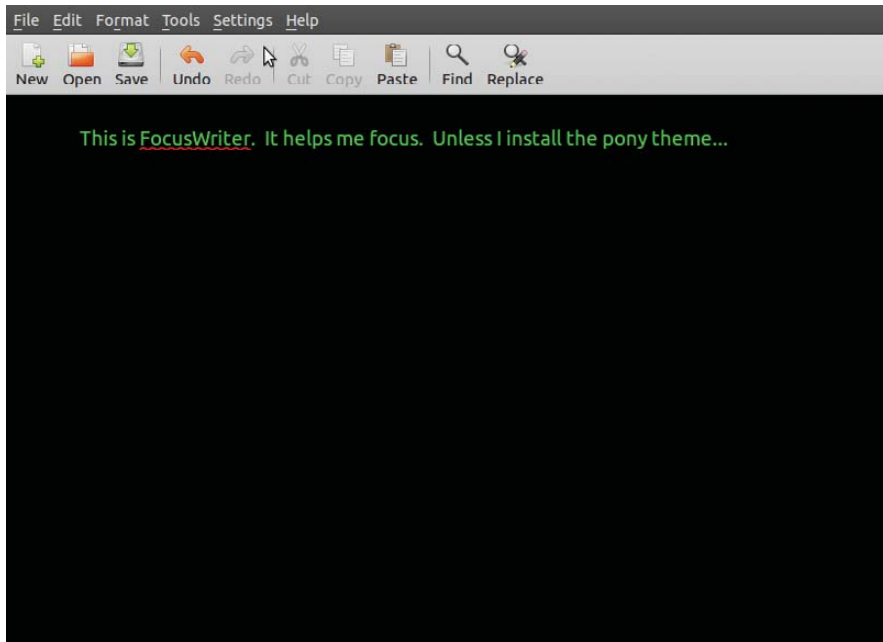


We use our staff.

visit our TS-7800 powered website at

www.embeddedARM.com

(480) 837-5200



Focus on Writing: FocusWriter

Because you're reading this in a magazine, it's fairly obvious I'm a writer. What might not be obvious, however, is that I'm also easily distracted—LOOK, A PONY! Seriously though, if you do any writing, you've likely discovered that word processors weren't designed for people that need to pump out a specific amount of text in a specific amount of time. There are a few features we don't need and really don't want:

- Formatting, smart quotes, rich text.
- Any format that is not plain text.
- Photo insertion.
- Automatic list formatting.

Those things are great if you're writing a paper for a college professor, or if you're creating a proposal for purchases. Unfortunately, what is a feature for some is a hindrance for others. As a writer, features I want are:

- A running word count—in publishing, word count is king.

- A full-screen mode to hide distractions—and ponies.
- Autosave.
- A standard and simple file/save interface.

And really, that's it. They may seem like simple things, but it's difficult to find software that meets the bill. A few options are available, but by far, the most usable I've found is FocusWriter. It has all the features I want, and it's small and fast requiring very few libraries to install. My only complaint is there's no option for keeping the menu on the screen. In order to see the menu, you need to mouse-over the top of the window. Yes, I know it's a feature, but it's a feature I'd like to turn off. The flickering menu tends to distract me. You know, like a pony.

Check it out at
gottcode.org/focuswriter.

—SHAWN POWERS

The MosKeyto's Buzz

A review of a USB drive might seem like a silly notion, but when the USB drive is barely bigger than the USB port itself, it seems worth mentioning. I recently was sent a LaCie MosKeyto USB drive, and I must admit, it's even smaller than I expected it would be. In fact, the cover to the Flash drive is actually bigger than the drive itself!

Because the drive is so small, it makes sense to keep it on a keychain. Most of my USB drives started with a cap, but after about three or four minutes of use, I would lose them. The LaCie drive has a neat keychain strap that secures both the drive and cap onto your keyring. So, if you're looking for a non-obtrusive USB drive to leave plugged in to your laptop, or if you want a portable Flash drive that won't dominate your keyring, the MosKeyto might be just the device you need. It also comes with some Wuala on-line storage, but really, what's exciting is how small these things are. They must compress the data or something—just teasing, it's obvious they fold the data to make it fit. Visit www.lacie.com/us/products/product.htm?pid=11546 for more info.



Associate Editor Shawn Powers and His MosKeyto

—SHAWN POWERS

SCALE 9X

The Ninth Annual
Southern California Linux Expo

Mark your calendars!


The 9th Annual
Southern California Linux Expo
is coming!

February 25-27, 2011

Expanded to five session tracks:
Beginner, Developer, SysAdmin
and two General Interest!

New, larger venue at the
Hilton Hotel @ LAX
Los Angeles, California

<http://www.socallinuxexpo.org>
Use Promo code LJAD for a 30% discount
on admission to SCALE

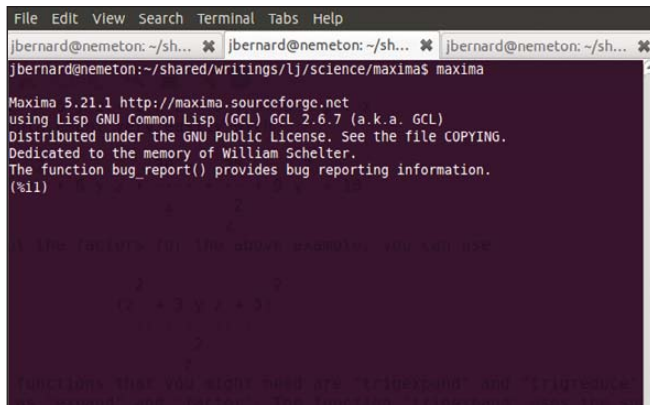


Maximum Theory with Maxima

Last month, I explained using Scilab to analyze the results of an experiment. In that specific case, I looked at a pendulum and how it behaved with different lengths and weights. This month, I cover using Maxima to help derive the equations describing the motion of a pendulum.

Maxima is a calculating engine that runs on the console, and several GUI front ends are available for it, such as wxMaxima. Maxima and its GUIs should be available for your distribution; if not, get the source code from the main Web site (maxima.sourceforge.com).

To start Maxima on the command line, type `maxima`. You should see some licensing information and then a prompt (Figure 1). To get a GUI, run `wxmaxima` (Figure 2).



```
File Edit View Search Terminal Tabs Help
jbernard@nemeton: ~/sh...  jbernard@nemeton: ~/sh...  jbernard@nemeton: ~/sh...
jbernard@nemeton:~/shared/writings/lj/science/maximas maxima
Maxima 5.21.1 http://maxima.sourceforge.net
using Lisp GNU Common Lisp (GCL) GCL 2.6.7 (a.k.a. GCL)
Distributed under the GNU Public License. See the file COPYING.
Dedicated to the memory of William Schelter.
The function bug_report() provides bug reporting information.
(%i1)
```

Figure 1. Maxima Prompt




Figure 2. wxMaxima GUI

Maxima runs as a continuous session, much like Maple or Mathematica. To end your session and exit the program, execute the `quit()` command. Notice that commands within Maxima end with a semicolon. Anyone used to coding in C should be familiar with this. If you end a command with a `$` instead, the output from that command won't be displayed.

Sometimes you may run a command in Maxima and not realize how long it will take to finish. If it's taking too long, stop execution by pressing `Ctrl-C`. Like in `bash`, `Ctrl-C` stops execution of the currently running command. Because this is a forced stop, you'll get an error

message, but don't worry about following up on it, unless you suspect something deeper is happening.

Maxima also contains help files for all of its built-in functions. To see detailed information about a function, type `? command` to see a help page on the given command. For most commands and functions, you'll also get some examples.

When you first start Maxima, the prompt is simply:

(%i1)

Every input line is labeled in a similar fashion, with incrementing numbers. So after the first prompt, the following input prompts would be `%i2`, `%i3` and so on. The output from commands are displayed immediately following the command, labeled with tags of the form `%o1`, `%o2`, `%o3` and so on. These behave the same as variables in Maxima. This means you can refer to earlier commands or their results simply by using the label. If you wanted to rerun the very first command you ran in this session, simply execute `%i1;`. There also are shortcuts for the last command and the last result. For the last command, simply type two single quotes (`' '`). For the last computed result, type a single percent sign (`%`). The percent sign also marks special built-in values—for example, `%e` (natural log base), `%i` (square root of -1) or `%pi` (3.14159...).

Maxima can handle arithmetic very easily. The most common arithmetic operations are:

- + — addition
- - — subtraction
- * — scalar multiplication
- / — division
- ^ or ** — exponentiation
- . — matrix multiplication
- `sqrt(x)` — square root of `x`

You can apply these to both numbers and variables. Arithmetic is done exactly, whenever possible. This means things like fractions are kept as fractions until you explicitly ask for a numeric result. So you see behavior like this:

```
(%i1) 1/100 + 1/101;
(%o2)          201
          -----
          10100
```

Variables are easy to use in Maxima. They simply are alphanumeric strings. Use the colon character to assign a value to a variable. For example, if you want to find the square of a number, use:

```
(%i1) x: 2;
(%o1)          2
(%i2) x^2;
(%o2)          4
```

You also can assign more complex functions to a variable—for example:

```
(%i3) w: sin(0.5 * %pi);
(%o3)          sin(0.5 %pi)
(%i4) w^2;
(%o4)          sin^2(0.5 %pi)
```

In certain cases, Maxima leaves results in a partly computed state. If you want to force Maxima to give you a number, add , numer to the end of your command. So, for the above, you would type:

```
(%i5) %o4, numer;
(%o5)          1
```

This use of numer will give you 16 significant figures, but Maxima can handle arbitrarily large numbers. The bfloat function converts your result to a numeric representation. The number of significant figures used by bfloat is set by the variable fpprec (the default is 16). As an example of really large numbers, try:

```
(%i1) %pi;
(%o1)          %pi
(%i2) %pi, numer;
(%o2)          3.141592653589793
(%i3) bfloat(%pi);
(%o3)          3.141592653589793b0
(%i4) fpprec: 100;
(%o4)          100
(%i5) bfloat(%pi);
(%o5) 3.1415926535897932384626433832795028841971693
    ➤9937510582097494459230781640628620899862
    ➤8034825342117068b0
```

Maxima handles algebra very easily. The expand function can be used to expand a polynomial—for example:

```
(%i1) (x + 3*y + z)^2;
(%o1)          (z + 3 y + x) 2
(%i2) expand(%);
()          2          2          2
          z + 6 y z + 2 x z + 9 y + 6 x y + x
```

You can do a variable replacement very easily with:

```
(%i3) %o2, x=5/z;
(%o3)          2          30 y 25          2
          z + 6 y z + ---- + -- + 9 y + 10
                      z          2
                      z
```

If you want to figure out the factors for the above example, use:

```
(%i4) factor(%);
(%o4)          2          2
          (z + 3 y z + 5)
          -----
                2
                z
```

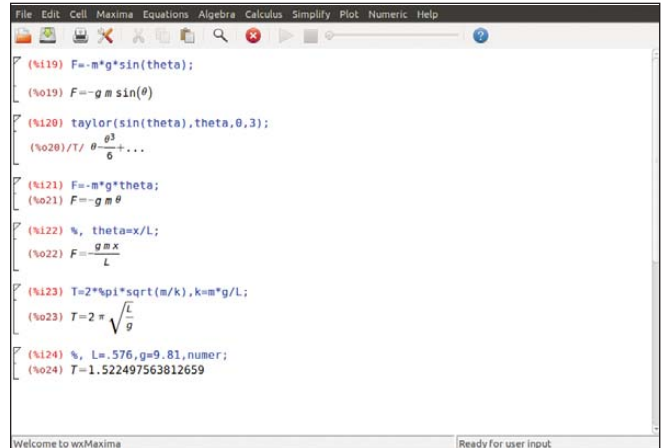


Figure 3. Time for an Oscillation

The last set of algebra functions that you might need are trigexpand and trigreduce. These do the same kind of function as expand and factor. The function trigexpand uses the sum of angles formulas to make the argument of each trigonometric function as simple as possible. trigreduce does the opposite and tries to make the expression such that it is a sum of terms with only one sin or cos function in each term.

Now, you should have enough information to use Maxima and figure out what you should have found last month with the pendulum. The only force that is applied to a pendulum is gravity, constantly trying to pull the pendulum straight down. This is called a restoring force and is:

$$F = -mg \sin(O)$$

where m is the mass of the pendulum bob, g is the acceleration due to gravity, and O is the angle the string makes with perpendicular pointing straight down. The sin function is a bit difficult to work with, so let's look at how to replace it. If you do a Taylor expansion of sin, you get:

```
(%i2) taylor(sin(x), [x], .0, 5);
(%o2)/T/          3      5
                  x      x
          x - --- + --- + . . .
                  6      120
```

From this, you can see that you probably can replace the sin function with O, as long as this angle isn't too big. This will give you:

$$F = -mgO$$

If the angle O is measured in radians, it's equal to x/L, where x is the actual amount of the arc that the pendulum bob travels along, and L is the length of the pendulum string. Once you do this, you've essentially minimized the nonlinear part to the point where you can treat the pendulum as a simple harmonic motion, with the force constant of (mg/L). If you plug this in to the equation for harmonic motion, you get the time for an oscillation given by 2*pi*sqrt(L/g). In a Maxima session, this would look like Figure 3. You can see that we get a result from theory that is really close to what we actually measured.

In the coming months, I'll look at more of Maxima's abilities, including advanced algebra, calculus and even plotting. Until then, go out and do some science of your own.

—JOEY BERNARD



REUVEN M. LERNER

Communication in HTML5

HTML5 gives Web applications new communication features.

Last month, I took an initial look at HTML5, the new standard for HTML that increasingly is supported by popular browsers. HTML5 includes a large number of different technologies and solutions, and it's being rolled out in a piecemeal fashion. Some parts of HTML, such as the abbreviated DOCTYPE declaration at the top of a document, can be used already. Other parts, such as new elements for HTML forms, are not quite ready yet for mainstream use, unless you are willing to use JavaScript to handle elements in browsers that lack support for those features.

For me, the improved form tags and expanded semantic markup for sections, headers and the like are both important and useful advances. But in many ways, I believe some of the most important improvements are a bit harder to understand and will take longer to catch on. These are improvements in how different pages may communicate with one another within the same browser or how the Web browser may communicate with outside resources.

To date, it has been difficult, if not impossible, for a program running in the browser (that is, with JavaScript) to create a mashup between two data sources or interact with another open window. This was because of a combination of historical, practical and security concerns, all of which were understandable. But today, we want our browsers to do more. Indeed, most modern applications are being developed for the browser, and if we can somehow push more information and intelligence to the browser, we'll both reduce the load on our servers and improve the responsiveness of the application.

A number of features in HTML5—and yes, some of these features aren't really part of the HTML5 specification, but I'll pretend they are for the purposes of this column—are designed to solve precisely this problem. These features aren't meant to make it easier to create Web pages, but rather Web applications. Specifically, I'm talking about interpage communication, WebSockets and threads, known as Web workers. Each of these topics probably deserves a column to itself, so I admit that the examples I provide here are meant to give you a taste of what's possible, rather than a comprehensive tutorial or example. Nevertheless, I hope you'll become excited about the possibilities raised here, and perhaps you'll even think of new and interesting ways to make use of these features.

Interpage Communication

The basic building block of the Web is the page, and more specifically, the page of HTML. Sure, modern Web browsers can display all sorts of different formats, but for all intents and purposes, when you talk about a Web app, you're talking about one or more pages of HTML. Nowadays, that page can be modified via JavaScript, using a variety of different techniques that often are lumped together with the Ajax name, regardless of how true or appropriate that is. Ajax has complicated things somewhat, removing the one-to-one correspondence that used to exist between HTTP requests and HTML pages. Although this means the flow of our applications has become more complicated, Ajax has made the Web a far more useful and friendly platform for users, and even for developers.

For many years, Web browsers were one-page-per-window affairs. If you wanted to browse three pages simultaneously, you needed to open multiple windows. Modern browsers all support the use of tabs, with some users (like me) abusing this feature quite a bit, opening dozens of tabs each day and then taking my sweet time to read and close the contents of each one. Given that all of these tabs (or windows) are running within the same program, it should be possible for them to communicate using a combination of JavaScript and HTML.

Perhaps this should be the case, but to date, it has not been possible. That's due to privacy concerns. You didn't want one Web page to be able to read from or write to another page without permission, and there wasn't a standard that would provide such permission. This is true not only in the case of two separate pages, but also in the case of two iframes on the same page, which might want to affect each other.

Now, if you're thinking you actually have been able to do this in the past without any hoopla, you might be right. It's true that an iframe can communicate with, and even modify, its parent window, but only if the two pages have the same origin. HTML5 changes the rules of the game by allowing pages to communicate with one another, even if they have different origins.

This works as follows. The sending page invokes the `postMessage()` method on the window or iframe that should receive the message, along with the expected origin of the receiver. For example, let's create a very simple HTML file that contains little more than an iframe (Listing 1). For now, ignore the

JavaScript event handler that is defined there. I'll get to that in a bit.

Inside this page is an iframe. Now, for demonstration purposes, this iframe will have the same origin as the outer page. But in many cases, the iframe will come from a completely different origin. In HTML5, that doesn't matter at all. You can send a message to whichever recipient you like. If you look at the HTML source for `iframe.html`, you'll see how to accomplish this:

```
$(document).ready(function() {
$("#send-button").click(function() {
    window.parent.postMessage($("#text-to-send").val(), '*');
});
});
```

In this example, I use jQuery to grab the button whose ID is "send-button". I then add an event handler to that button, indicating that when it is clicked, it should invoke `window.parent.postMessage`, sending the text contained inside the text field. I should note that the `postMessage()` method can be called on any window or iframe, and that it can send any text in its first parameter.

The second parameter indicates the origin of the recipient to whom you're sending this message. In this case, I have indicated that the recipient may have any origin by specifying a wild card. In production environments, it's probably safe to assume you will want to specify an origin. By stating the recipient's origin, there's a bit of additional safety—the message will be sent only if the receiving window object's content is from the stated origin.

On the receiver's end, the posted message arrives as an event, one which the receiver can (and should!) examine before using. Going back to `atf.html`, you will see how the receiver accepts a message in its event handler:

```
$(document).ready(function() {
window.addEventListener('message', receiver,
false);
function receiver(e) {
    alert("origin = " + e.origin + "");
    alert("data = " + e.data + "");
    $("#message").text(e.data);
};
});
```

The event handler for this page indicates that it's willing to accept a message. Each message consists of two pieces, the message (the text string that the sender passed as a parameter to `postMessage`) and the origin (the sender's origin). Note that the sender cannot set its origin; this piece of information is handled by the browser.

Listing 1. `atf.html`, the Container HTML Page

```
<!DOCTYPE html>
<head>
<title>Page title</title>
<script src="jquery.js"></script>

<script>
    $(document).ready(function() {
        window.addEventListener('message', receiver, false);
        function receiver(e) {
            alert("origin = " + e.origin + "");
            alert("data = " + e.data + "");
            $("#message").val(e.data);
        };
    });
</script>

</head>
<body>
<h1>Page headline</h1>
<p id="message">[No message yet]</p>

<iframe id="my-iframe" src="iframe.html" />
</body>
</html>
```

Listing 2. `iframe.html`, the Contents of Which Are Loaded into an iframe

```
<!DOCTYPE html>
<head>
<title>iframe title</title>
<script src="jquery.js" /></script>
</head>
<body>
<h1>iframe headline</h1>

<p>Text to send: <input type="text" id="text-to-send" /></p>
<p><input id="send-button" type="button" value="Send it!" /></p>

<script>
$(document).ready(function() {
    $("#send-button").click(function() {
        window.parent.postMessage($("#text-to-send").val(), '*');
    });
});
</script>

</body>
</html>
```

Because the origin information is passed along with the message, it's possible for the receiver to filter out which origins it is willing to accept. In other words, although it's possible a rogue site will try to start posting to random windows that you might have open on

other sites, the only way such messages actually will arrive is if the receivers are willing to accept them. I'm sure someone with more of a black-hat mentality than mine will find ways to defeat this security mechanism, but from what I can tell, it was thought out very carefully and cleverly, and should avoid most mischief.

Now that it's possible for any window to send messages to any other window, what can you do with it? The answer, of course, is that no one knows. Off of the top of my head, I can imagine chat clients—or more generally, using a single window on a Web browser as a communication switchboard and clearing-house—grabbing feeds and incoming messages and putting them on the appropriate pages. Imagine if Facebook were to have a single iframe that would handle its (very large!) number of interactions with the server, and then handle all page updates through that iframe, rather than on each individual window or tab.

I also can imagine the `postMessage()` method ushering in a new age of multiwindow, desktop-like applications. Think of how many desktop applications now use multiple windows—one for control, another for each document and yet another for a “palette” of options. Now you can do the same thing with a Web browser, with a native message-passing interface.

Just what people will do with these capabilities is unknown, but I predict we'll see a rash of new, rich, browser-based applications that take advantage of them.

WebSockets

One of the greatest contributions of UCB (Berkeley) to the UNIX operating system was the introduction of sockets. Sockets allowed programmers to open a connection to another computer easily and quickly. Once opened, the socket operated something like a point-to-point file handle, allowing you to ignore the fact that data in the socket was being transmitted through dozens or hundreds of other computers. A large number of Internet services, from SMTP to FTP to HTTP, use sockets. I personally have used them for nearly 20 years to implement everything from my undergraduate thesis, to Web browsers and servers, to various Internet-enabled applications.

HTML5 brings socket-like connections to the browser, using a technology called WebSockets. WebSockets are similar in principle to UNIX sockets, in that you can open a connection to an arbitrary other point on the Internet, and send and receive data reliably without even considering the numerous hops or connections along the way.

Now, if you are an experienced Web developer, you might be wondering what the big deal is. After all, Ajax calls allow you to open HTTP connections and send and receive data. And `xhr` (the `XmlHttpRequest` function) has been around for a few years, and it works quite well. The difference is WebSockets will allow you to open one or more connections to

Listing 3. ws.html

```
<!DOCTYPE html>
<head>
<title>Page title</title>
<script src="jquery.js"></script>

<script>
  var weatherSocket = new WebSocket("ws://localhost:8080");

  $(document).ready(function() {

    weatherSocket.onopen = function(e) {
      alert("Opened weather socket");
    };

    weatherSocket.onmessage = function(e) {
      alert("Received a message: " + e.data);
    };

    weatherSocket.onclose = function(e) {
      alert("Closing the weather socket...");
    };

  });
</script>

</head>
<body>
<h1>Page headline</h1>
<p>WebSockets!</p>

<script>

  while(weatherSocket.readyState == 0)
  {
    alert("socket state is " + weatherSocket.readyState);
  }

  alert("socket state is " + weatherSocket.readyState);

  weatherSocket.send("Hello from the client!");
  alert("socket state is " + weatherSocket.readyState);
  weatherSocket.close();
</script>

</body>
</html>
```


anywhere on the Internet, not just to servers with the same origin as the current page. Moreover, WebSockets use their own protocol that is admittedly quite similar to HTTP, but it has a great deal less overhead. Finally, WebSockets remain open as long as the sides agree to do so—as opposed to HTTP, which is meant to be stateless and to be closed after a single request-response transaction. For all these reasons, communication using WebSockets generally is going to be far more efficient. A number of articles describing WebSockets have done the math and show just how much more efficient WebSockets are than HTTP—and the difference is staggering.

Working with WebSockets is remarkably simple. You open a WebSocket with some JavaScript code, often (presumably) fired either when a user performs an action (such as pressing a button) or when a certain event takes place (for example, a certain amount of time has elapsed). No matter what, you open a new WebSocket by specifying the URI to which you want to connect, starting with a protocol name (either ws or wss, for unencrypted or encrypted, respectively), continuing with the hostname, and then ending with a resource name.

Once the WebSocket is open, you can attach callbacks to it, indicating what should happen when the socket is opened, closed or receives a message. (Each time the WebSocket receives data from the remote host, it will invoke the “onmessage” callback function.)

For example, here’s a simple WebSocket that retrieves data from a hypothetical weather server:

```
var weatherSocket = new WebSocket("ws://localhost:8080");  
    // Our own weather server
```

Then, you can assign callbacks:

```
weatherSocket.onopen = function(e) {  
    alert("Opened weather socket");  
};  
  
weatherSocket.onmessage = function(e) {  
    alert("Received a message: " + e.data);  
};  
  
weatherSocket.onclose = function(e) {  
    alert("Closing the weather socket...");  
};
```

Small, Portable Devices with Ubuntu Linux

Small Form Factor Intel® Atom™ Platform

No fans, no moving parts. Just quiet, reliable operation. Incredibly compact and full featured; no compromises.



VESA-Mountable NVIDIA® ION/ION2 System

Compact, lightweight system with GeForce® Graphics. Flexible storage options (dual HDD support) and WiFi.

Value only an **Industry Leader** can provide.

Selecting a complete, dedicated platform from Logic Supply is simple: Pre-configured systems perfect for both business & desktop use, Linux development services for greater system customization, and a wealth of online resources all within a few clicks.

[Learn More > www.logicsupply.com/linux](http://www.logicsupply.com/linux)

LOGIC
SUPPLY

Finally, you can send messages by invoking the `send()` method—yes, the same method that you saw above, but without the second parameter indicating the origin.

Notice that although you write directly to the `WebSocket` using `send`, you don't read a result directly from it or via a return value to `send()`. Rather, you will get the data when it is sent to you, via the execution of your method at `weatherSocket.onmessage()`.

One piece is missing from this description, namely a server to which the `WebSocket` connects. You cannot connect to just any old server on the other end, and especially not to an HTTP server. Fortunately, a growing number of packages (in various open-source languages) can handle the server side of `WebSockets`. One such package is the `em-websocket` gem for Ruby, based on the well-known `eventmachine` gem. `WebSocket` server libraries already exist for PHP and Python, as well as a number of other languages. Over time, I expect to see a number of `WebSocket`-compatible servers emerge.

How can you use `WebSockets`? As with interwindow communication, I expect the best applications and ideas haven't been developed yet. But once your Web browser can connect to any host on the Internet using a specialized high-performance protocol, you can

Indeed, Web workers operate almost as if they existed on different computers, with no direct connections between them.

Imagine that the sky is the limit. Suddenly, Web-based chat servers no longer need to use kludges or hacks in order to allow for real-time chat. You can create mashups on the client, rather than the server. Combined with the new geolocation facilities in HTML5, you can have a map that updates your location in real time, using nothing more than HTML and JavaScript. It does mean that on the server side, Web applications now will require more than just installing Apache, but that has been true for a while now, as applications have become more sophisticated, so I don't think you need to worry about that too much.

Web Workers

Finally, another intriguing addition to HTML5 is the notion of Web workers, which you can think of as the browser equivalent of threads. Perhaps you have a complex task that needs to be handled in parallel with rendering of a page or of downloading information from the Internet. By splitting the work across two Web workers, you can take advantage of today's modern, multiprocessor computers to get a faster result. Because Web workers operate in the background, rather than on the thread that handles displaying the

page, the page should be more responsive than if everything were in one thread.

Now, I must admit I generally have tried to avoid programming with threads, because of the many problems that can crop up whenever you have shared resources. Given that JavaScript was never designed to work with threads, my first thought when I heard about Web workers was how this possibly could work while keeping data safe and the browser stable. The solution appears to be sound, although it's still too early for me (and many others) to know for sure.

The idea is this. You launch a Web worker by creating a new `Worker` object:

```
var worker = new Worker("code.js");
```

Notice that you hand the name of a file to a Web worker. You cannot hand it a piece of code, either directly or by passing it a function reference. Perhaps it eventually will be possible for a browser-based application to create dynamically, and then store (using `WebSockets`?) a file on the server, but the main purpose of this restriction is to ensure that there is no chance for shared data among the various Web workers, thus avoiding the chance for issues traditionally associated with threads.

Indeed, Web workers operate almost as if they existed on different computers, with no direct connections between them. Workers cannot access the DOM, which means any elements on the page. Functions and data in the main thread are not available to the Web workers, and vice versa.

This raises the question of how the main thread and Web workers can communicate. The answer probably won't surprise you. They use `postMessage()`, the same mechanism for message passing that can be used to send information from one window or tab to another, regardless of origin.

I can foresee a number of uses for Web workers. First, they will allow browser-based applications to handle more than one thing at a time, ensuring that the main thread is used for rendering the UI and interacting with the user. Second, it means you can start to break problems apart, taking advantage of modern computer hardware that can put different threads on different processors intelligently. Finally, it means JavaScript now has the beginnings of a built-in message-passing mechanism. And, although programs still must remain inside a single browser, I have to assume at some point, it'll be possible to open a Web worker not just on your local machine, but on a remote one, as well.

Conclusion

Marc Andreessen, who wrote the original Mosaic browser and helped to popularize the Web as a founder of Netscape, claimed years ago that the browser is the new operating system. Even as Ajax and other advanced Web

technologies have advanced during the past few years, and such amazing browser-based applications as Google Docs have emerged, I still have been skeptical of whether Web-based applications ever will truly rival their desktop counterparts. The addition of cross-window communication, WebSockets and Web workers go a long way toward convincing me that Andreessen's prediction has nearly come true.

HTML5 and its associated technologies include a wealth of new options for developers. It will take some time to figure out how well these work, how to get around the fact that not all browsers support them and just how useful (or not) various features might be. If you are a Web developer, I encourage you to study and work with these technologies as soon as possible. I already have changed the architecture of some of my applications as a result, and I wouldn't be surprised if that happens to you too. ■

Reuven M. Lerner is a longtime Web developer, architect and trainer. He is a PhD candidate in learning sciences at Northwestern University, researching the design and analysis of collaborative on-line communities. Reuven lives with his wife and three children in Modi'in, Israel.

Resources

The best book I've read on the subject, *Dive Into HTML5*, isn't even a proper book at the time of this writing, but rather a free on-line resource written by Mark Pilgrim (diveintohtml5.org). If you are familiar with Pilgrim's previous work, such as *Dive Into Python*, you know his writing is excellent. Not surprisingly, this was the first resource to which I turned to bone up on HTML5, and it continues to be my favorite combination of tutorial and reference.

However, Pilgrim's book says very little on the subjects I've mentioned in this month's column. For excellent tutorials on these subjects, I recommend *Pro HTML5 Programming* written by Peter Lubbers, Brian Albers and Frank Salim. This last book also is aimed at beginners. Although I think other books are better than this one in other areas, it really shines in the cases I mentioned this month.

Finally, take a look at www.html5rocks.com, a Google-sponsored site that describes various HTML5-related technologies, including documentation and code examples.



visit us at www.siliconmechanics.com
or call us toll free at 866-352-1173



As the Chief Financial Officer for Silicon Mechanics, Steve is an Expert where value is concerned. That's why he's pictured here with the Rackform iServ R143.

The R143 is a flexible and affordable 1U server. It features an Intel® Xeon® Processor 3400 Series, with powerful features like Turbo Boost, Hyper-Threading, and DDR3 memory. This processor is also available in a low-voltage version, which can optimize power usage and help contain energy costs. With 6 DDR3 DIMM sockets, 2 Gigabit Ethernet adapters, a PCIe expansion slot, and 4 hot-swap SAS/SATA drive bays, the R143 can handle a lot more than entry-level workloads. With a price that starts around \$1250, you don't have to be a CFO to understand the value.

When you partner with Silicon Mechanics, you get more than a flexible, affordable entry-level server — you get an Expert like Steve.

For more information about the
Rackform iServ R143
visit www.siliconmechanics.com/R143

Expert included.

Silicon Mechanics and the Silicon Mechanics logo are registered trademarks of Silicon Mechanics, Inc. Intel, the Intel logo, Xeon, and Xeon Inside, are trademarks or registered trademarks of Intel Corporation in the US and other countries.



DAVE TAYLOR

Dealing with Spaces in Filenames

How do you deal with this problem in your scripts?

In the good old days when UNIX was young, no one ever would have thought of putting a space in a filename. It simply wasn't done—just as you'd never do that on a DOS or Windows system. Filenames were short, succinct and well-formed, like HW43.DOC.

Most of the Linux command-line utilities and the shells themselves have been designed based on the premise that a space delimits a field value rather than being an acceptable component of a filename. If you've done any scripting, you already know this. Spaces in filenames can cause great trouble in shell scripts! Here's a simple example:

```
for name in $(ls | grep a)
do
  echo "File #$count = $name"
  count=$(( $count + 1 ))
done
```

To set the stage, I've created a directory with some tricky filenames:

```
$ ls
"quoted" beastly filename      sample2.txt
multi-word file name.pdf      test.sh
```

Yes, to maximize trouble, I have a filename that includes quotes as well as a space. Don't get me started on having an escape character or non-printable character in the name though. It's doable, but I'd rename it as soon as possible.

Not all the filenames above have an "a" in them, so let's see what happens when the fragmentary script is run in this directory:

```
$ ./test.sh
File # = "quoted"
File #1 = beastly
File #2 = filename
File #3 = multi-word
File #4 = file
File #5 = name.pdf
File #6 = sample2.txt
```

Oh, is that ugly and wrong!

The shell can deal with these filenames if they're simple enough, and the for loop for `name in *` yields three filenames, not six, but somewhere along

your scripting journey, you inevitably will slam into the problem of embedded spaces.

The most common error is to forget to quote filenames when you use them elsewhere in the script, of course. As an example, let's work on a script that replaces spaces in filenames with underscores.

File Renaming with a Bug

The obvious solution to such renaming is something like this:

```
for name in "*" *
do
  newname="$(echo $name | sed 's/ /_/g')"
  mv $name $newname
done
```

This doesn't work, however, and in a most fascinating way:

```
mv "quoted" beastly filename multi-word file
  ↳name.pdf sample2.txt test.sh "quoted"
  ↳beastly filename multi-word file
  ↳name.pdf sample2.txt test.sh
↳"quoted"_beastly_filename_multi-word_file_
↳name.pdf_sample2.txt_test.sh"quoted"_beastly_
↳filename_multi-word_file_name.pdf_sample2.txt_test.sh
```

What's happened is that `"* *"` simply produces two full filename listings rather than just those filenames that contain a space—oops. Let's try a different pattern:

```
for name in *\ *

```

That does the trick, but we've not compensated for the fact that when the shell sees a line like:

```
mv multi-word file name.pdf multi-word_file_name.pdf
```

it's going to complain that it's seeing four filename arguments to the `mv` command, not the required two:

```
usage: mv [-f | -i | -n] [-v] source target
        mv [-f | -i | -n] [-v] source ... directory
```

In this case, the solution is to quote the filename variable:

```
mv "$name" $newname
```

As a discipline, it's always good to quote filenames you reference in any context to ensure that when the shell passes them to the command as arguments, the filenames with embedded spaces are handled properly.

This isn't a universal solution, however, because if you're using subshells and pipes, it can be pretty darn hard for quotes to survive multiple steps.

One path to travel is to set IFS, the internal field separator, in the shell to something other than a space, as explained in the Bash man page:

IFS: The Internal Field Separator that is used for word splitting after expansion and to split lines into words with the read built-in command. The default value is "<space><tab><new-line>".

That's useful for "read", particularly if you're reading lines of text and want to have a different field separator (think flat-file text database files), but it still doesn't really solve our filename problem.

One thing I've used in the past, although it's a sloppy, crude solution, is to start by changing spaces to some unlikely sequence of characters, run through all the processing, and change them back at the last second. For example:

```
safename="$(echo name | sed 's/ /_/_/g')"
```

and reversed with:

```
original="$(echo $safename | sed 's/_/_/g')"
```

It solves the problem, but it's definitely not a very efficient or smart use of computing resources.

I've outlined three possible solution paths herein: modifying the IFS value, ensuring that you always quote filenames where referenced, and rewriting filenames internally to replace spaces with unlikely character sequences, reversing it on your way out of the script.

By the way, have you ever tried using the find|xargs pair with filenames that have spaces? It's sufficiently complicated that modern versions of these two commands have special arguments to denote that spaces might appear as part of the filenames: find -print and xargs -0 (and typically, they're not the same flags, but that's another story).

During the years I've been writing this column, I've more than once tripped up on this particular problem and received e-mail messages from readers sharing how a sample script tosses up its bits when a file with a space in its name appears. They're right.

My defensive reaction is "dude, don't use spaces in filenames", but that's not really a good long-term solution, is it?

What I'd like to do instead is open this up for discussion on the *Linux Journal* discussion boards: how do you solve this problem within your scripts? Or, do you just religiously avoid using spaces in your filenames? ■

Dave Taylor has been hacking shell scripts for a really long time, 30 years. He's the author of the popular *Wicked Cool Shell Scripts* and can be found on Twitter as @DaveTaylor and more generally at www.DaveTaylorOnline.com.



Linux - FreeBSD - OpenSolaris - etc.

Proven Technology.

Proven Reliability.

When you can't afford to take chances with your business data or productivity, rely on a Genstor server customized to your specifications.

POWER

PERFORMANCE

Fly into the Cloud with Genstor Systems



- Up to 48 cores in a 1U.
- AMD Opteron 6100 series.
- Single high-efficiency power supply.
- Up to 512GB DDR3 memory.
- Ideal as front end processing servers.



- Up to 12 cores in a 2U
- Dual redundant high efficiency power.
- Up to 96GB DDR3 memory.
- Server Power Capping via Intel Intelligent Power Node Manager.
- Ideal as front end processing and/or storage.



- Up to 4 GPU cards.
- Dual redundant high efficiency power.
- Up to 192GB DDR3 memory
- Up to 24 cores using 2 CPUs.
- Up to 8 3.5" disks.

Genstor Systems, Inc.



Powerful.
Intelligent.

1501 Space Park Drive
Santa Clara, CA 95054
www.genstor.com

E-mail: sales@genstor.com
Phone: 877-25 SERVER
408-980-0121

Intel®, the Intel® Logo, Intel® Xeon®, and Xeon® Inside® are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.



KYLE RANKIN

Status Messages in Screen

If you live in screen, learn how to make its status bar notify you like a pop-up window on a desktop.

If you have used the command line for any length of time, someone probably has introduced you to a program named `screen`. If you aren't familiar with it, `screen` allows you to manage multiple console sessions within the same program, so instead of dealing with tabs or multiple terminal windows, you can shift between your console sessions with a few key-presses. In addition to this, `screen` allows you to detach from these console sessions completely while they are still running and then reattach to them later. One use for this could be to start a process that might take some time near the end of the day at work within a `screen` session, detach from the session, then reattach later on that evening at home to make sure it completed.

I use `screen` along with `lrssi` on a server I leave running, so I can stay logged in to IRC at all times and

Special string escapes also allow you to specify the current time, the hostname and system load.

just attach to that `screen` session from any machine I happen to be using. I also keep other console sessions in `screen` set up for `mutt` and SSH instances into various servers, so I can access everything from the same session. To be honest, I spend a majority of time in front of a computer within `screen` sessions.

Because I devote so much of my focus to a `screen` window, I've found it handy to take advantage of the status line within `screen` to notify me of anything that might need my attention. The status line often is off by default; however, once enabled, it takes up a row at the bottom of your console. You can configure it to list all sorts of information, from things as basic as the current date and time to as complicated as the output of any shell program you can imagine.

Configuring the status line in `screen` quickly can become quite complicated. There is a robust syntax of string escapes that allows you to control right and left justification and coloration of all the items in the status line. Special string escapes also allow you to specify the current time, the hostname and system load. Describing all these options is much like describing `printf` syntax to someone for the first time. Instead of

going into all of that here, I refer you to the section of the `screen` man page titled STRING ESCAPES. To get there, type `man screen`, and once you are in the man page, type `/^STRING ESCAPES`, and press Enter.

For this column, I show how to enable the status line in `screen` and tweak it so it shows you some basic system information. Finally, I explain how to add custom script output. If you decide you want to have fancy colorized status lines, I'm sure you'll have fun experimenting with all of the options in the man page.

Enable the Status Line

Each user's individual `screen` settings are configured in `~/.screenrc`, so to enable a basic status line, you use the `hardstatus` configuration option. To create a status line that says "hello world", add the following lines to your `~/.screenrc`:

```
hardstatus alwayslastline
hardstatus string 'hello world'
```

Now when you start `screen`, you should see "hello world" along the bottom of your window in reverse video (that is, your foreground and background colors will be reversed). Some people like that, but if you want to have white text instead, change the `hardstatus` string to:

```
hardstatus string '%{= w}hello world'
```

The next time you load your `screen` session, the text will be white. If you want, you also can change the `.screenrc` from within the `screen` session itself and reload it. Once you have saved any changes to your `.screenrc`, just press `Ctrl-a :` (that's `Ctrl-a` followed by the `:` key), then type `source ~/.screenrc`, and press Enter.

Since the status line works now, let's make it display more useful information, like the current date and time, and system load. Change your `hardstatus` string to:

```
hardstatus string '%{= w}%Y-%m-%d %c | %l'
```

This looks a bit complicated, but let's break it down. The `%{= w}` part of the `screen` sets the foreground to white. I can represent the year, month and day with `%Y`, `%m` and `%d`, respectively, and in this case,

I added a hyphen (-) in between each value so it was easier to read. Then, I added a space and %c, which is expanded into the current time. Finally, I added a | symbol with some spaces for padding and %l, which expands out into the current load.

Custom Status Scripts

As you can see in my example, screen provides a few string escapes for some common output you might want in the status line. It also allows you to define custom commands it can run and display the output on the status line instead. All you have to do is define a command in your .screenrc via the backtick option, and then reference that command in the hardstatus string. For instance, here's a simple Perl script I wrote that parses the output of `fetchmail -c` (which checks an IMAP account of mine for new messages). If any of my defined folders have new messages, it outputs them separated by spaces. Name the following script `/usr/local/bin/new_mail_check.pl`:

```
#!/usr/bin/perl

open FETCHMAIL, "/usr/bin/fetchmail -t 10 -c 2>/dev/null |"
  or die "Can't run fetchmail: $!";

while(<FETCHMAIL>){
    if(/^(\d+) messages \((\d+) seen.*?folder (.*)\)/){
        $m+=$1; $s+=$2; $f=$3;
    }
    # you might have to change this regex depending on
    # how your IMAP server displays subfolders
    $f =~ s/INBOX\././;
}
if($1 > $2){
    $fs{$f} = $1 - $2;
}
}
close FETCHMAIL;
$t = $m - $s;

if($t > 0){
    foreach $folder (sort { $fs{$a}<=>$fs{$b} } keys %fs){
        push @folders, "$folder:$fs{$folder}";
    }
}
print join " ", @folders;
```

Make sure the script is executable, then change your `~/.screenrc` to the following:

```
backtick 101 60 60 /usr/local/bin/new_mail_check.pl
hardstatus alwayslastline
hardstatus string '%(= w)%Y-%m-%d %c | %l | %101`'
```

The first line defines a backtick command that will be referred to as 101. The two 60s that follow

```
Welcome to the social channel for Linux Journal - The Original Magazine of the Linux Comm
[!:#linuxjournal(+Centz)][@greenfly(+i)][Act: 2,5,6,7,9]
19:40 biosshadow| its good
19:40 biosshadow| just not cold
19:40 KitsuneDrag0n| biosshadow: build a fire out of your roommate
19:40 the_bdquick| so LaRue is using the TRex defense on bpx-dev
19:41 LaRue| TRex?
19:42 the_bdquick| staying where you are, so sight is based on movement
19:42 KitsuneDrag0n| LaRue: are you a dinosaur?
19:42 LaRue| rawr?
19:42 KitsuneDrag0n| :D
19:43 biosshadow| KitsuneDrag0n yea no
19:43 KitsuneDrag0n| i would want to be a stegosaurus
19:43 KitsuneDrag0n| biosshadow: you DONT want to put lube in the microwave???
19:48 jrabbit| [-babyteal@unaaffiliated/jrabbit] has quit from #linuxjournal [Read
error: Connection reset by peer]
19:48 dbart| [-dbart@cpe-174-109-078-104.nc.res.rr.com] has quit from
#linuxjournal [Read error: Connection reset by peer]
19:48 mode/#linuxjournal [+o dbart] by ChanServ
19:48 mode/#linuxjournal [-dbart@cpe-174-109-078-104.nc.res.rr.com] has joined #linuxjournal
19:52 jrabbit| [-babyteal@unaaffiliated/jrabbit] has joined #linuxjournal
19:58 Judas_PhD| [-kevin@misterfluffy.dsl.xmission.com] has joined #linuxjournal
20:04 Judas_PhD| [-kevin@misterfluffy.dsl.xmission.com] has quit from #linuxjournal
[Quit: This is a quitting message]
20:07
2010-11-02 20:07| |0.19|INBOX:1 yammer:1|19:51:49 shawnp0wers
```

Figure 1. A Sample of My Decked-Out Screen Status

define the life span and auto-refresh times for the command in seconds. The lifespan is defined as the number of seconds the output is considered valid before the command will be run again if the string escape is encountered. The auto-refresh value defines when to refresh the display of the hardstatus string. I usually define both values to be the same for my backtick commands. In this case, I check for new mail every 60 seconds. The final argument in the command is the full path to the command you want to run. If you need to specify any arguments, you can put them after the command. Also notice that I added %101` to the end of my hardstatus string. The %` string escape will put the specified backtick output into the status. In this case, %101` will put the output of the backtick command I define as 101. If I wanted to add another command, I'd add another backtick line to my .screenrc and define it as 102.

So, where do you go from here? Well, besides using all the string escapes to colorize your status lines, you really are limited only by your shell scripting ability. For instance, there's a program for Irssi called `fnotify` you can use to log everyone who highlights your handle in Irssi into a file. I wrote a simple script to pull the last line from that file, parse it, and display the date and handle of the last person who talked to me on IRC to my status bar. You also could write scripts to show you information from Twitter, headlines from your favorite sites, Nagios alert summaries or even random output from the `fortune` command. If you are like me though, you'll spend the bulk of your time tweaking all the colors and left and right padding, so your status looks just right. ■

Kyle Rankin is a Systems Architect in the San Francisco Bay Area and the author of a number of books, including *The Official Ubuntu Server Book*, *Knoppix Hacks* and *Ubuntu Hacks*. He is currently the president of the North Bay Linux Users' Group.

LSI's Engenio 2600-HD High-Density Storage System

If you're using the Lustre FS on your HPC system, you might be able to improve your performance with LSI Corporation's new Engenio 2600-HD, a high-density storage system that delivers a reported 40GB/s of throughput and scaling to 1.8 PB of capacity in a single standard rack. LSI says that Engenio's 2600-HD's highly scalable, dense architecture helps HPC organizations maximize productivity and achieve a quicker time to results, while minimizing data-center floor space and overall energy consumption. The system consists of two LSI 6Gb/s SAS-based controllers integrated into the new Engenio DE6600 high-density SAS drive enclosure. The system is capable of sustaining up to 4GB/s of throughput and housing up to 60 SAS drives in a 4U space.

lsi.com/hpc



Opengear's ACM5004-G Mobile 3G Cellular Router

The latest open gear from Opengear is the company's new ACM5004-G mobile 3G cellular router for secure high-speed wireless connectivity to remote sites and devices. The compact, industrial-grade device, which delivers real-time access, monitoring and control regardless of location, has an open-source Linux core and offers local custom scripting. Key features include ubiquitous routing, secure remote control, extensive monitoring and alerts, remote power management, support for custom apps and external USB.

www.opengear.com



Wolfram Research's Mathematica



Mathematica from Wolfram Research, a favorite tool of Linux-geek number-crunchers everywhere, recently crossed the magic threshold of version 8. The new version 8 of this powerful computation, development and deployment platform adds free-form linguistic input via its novel Wolfram|Alpha technology, which enables users to input plain English and get immediate results without the need for syntax. Among the other 500 new additions are improved capabilities for statistical distributions and data visualization, built-in GPU programming support, SymbolicC support, integrated control systems, wavelets functions, option pricing solvers and feature detection in image processing. Mathematica 8 is available for Linux x86, Mac OS X and Windows XP/Vista/7.

www.wolfram.com/mathematica

```

plot x^2+y^2<2 and x^3-y^3<1 with dashed grid lines
Result
Show[RegionPlot[x^2 + y^2 < 2 && x^3 < 1 + y^3,
{x, -2.2, 2.2}, {y, -2.3, 2.1},
{GridLines -> Automatic, GridLinesStyle -> Dashed}]]

```



Martin Hüdepohl's *Badass LEGO Guns* (No Starch Press)

With a title like *Badass LEGO Guns*, how can you not judge a book by its cover? This fun new book by Martin Hüdepohl and published by No Starch Press illustrates how to build five eclectic weapons entirely from LEGO Technic parts that can shoot plastic LEGO bricks at high speed with a high level of accuracy. The builder adds only rubber bands, some sanding and a touch of Krazy Glue to build these functional *fusils*, each with its own kick-butt nickname: the Warbeast submachine gun, the Thriller and Mini-Thriller crossbows, the Parabella mini-marvel and the Lilliputt semi-automatic pistol with a nine-brick magazine. The models range from sophisticated to simple, and "builders of all ages will find something enjoyable", notes the publisher.

www.nostarch.com



R. B. Thompson and B. F. Thompson's *Building the Perfect PC*, 3rd ed. (O'Reilly Media)

Robert Bruce Thompson and Barbara Fritchman Thompson, authors of the new 3rd edition of *Building the Perfect PC*, say you don't even need to be a geek to build your own PC. Well, we are geeks and we want to build our own PCs too. As talented as we are though, we may want to pick up the Thompson team's updated book to make sure we don't blow it. The payoff is a PC that is of higher quality and lower cost than off-the-shelf models. The authors explain what components you'll need as well as where to find them. They also explain how to build for your OS of choice and take advantage of the latest multicore CPUs. Instructions cover how to build numerous types of PCs, including a general-purpose computer, an extreme gaming machine, a media center, an appliance, a low-cost PC or a home server.

www.oreilly.com

TYAN's AMD FireStream GPU Compute Accelerators

Four new lines of server platform are now available from TYAN, all of which are designed to take full advantage of AMD FireStream GPU compute accelerators. These compute accelerators deliver what TYAN labels "shocking floating-point performance that is exponentially faster than x86 CPUs in some applications". The solutions—models B7015, S7025, S8225 and S8236—range from one AMD FireStream compute accelerator in a 1U server up to eight compute accelerators in a 4U platform. These platforms feature double-wide PCIe 2.0 x16 slots, and they meet the special mechanical requirements as well as the power and airflow needs to support AMD FireStream 9170, 9250, 9270, 9350 and 9370 GPU compute accelerators.

www.tyan.com



QA Graphics' Energy Efficiency Education Dashboard

As the old adage goes, you can't manage what you don't measure. And if managing resource (such as electricity, water and so on) use or production is your goal, QA Graphics' Energy Efficiency Education Dashboard (EEED), just upgraded to version 2.0, may be the tool for the job. EEED is an interactive solution that displays real-time building data and educates occupants on sustainable building-management practices. The new version is a fully encapsulated application based on the Adobe AIR platform that runs as a standalone client or software as a service with no browser constraints. The solution often is used by organizations to help earn points toward

green-building-certification programs, such as US Green Building Council's LEED and the Green Building Initiative's Green Globes.

www.qagraphics.com/eed

SuperLumin Networks' Nemesis

Don't got bandwidth? Don't get bummed; get SuperLumin Networks' Nemesis, a 64-bit caching and application-acceleration platform. Besides offering the capabilities of a standard proxy cache, Nemesis addresses the need to enable and cache streaming video and rich-media on the Web. The application is optimized to cache bandwidth-intensive social-media sites, such as Facebook and YouTube. Other features include scalability to 64 CPUs, 100 million cache objects, policy-based content filtering, bandwidth management/traffic shaping and content distribution, among others. Nemesis runs on SUSE Linux.

www.superlumin.com



Please send information about releases of Linux-related products to newproducts@linuxjournal.com or New Products c/o Linux Journal, PO Box 980985, Houston, TX 77098. Submissions are edited for length and content.

Fresh from the Labs

downtimed—System Downtime Monitor

dist.epipe.com/downtimed

The following statement appears at the top of the *downtimed* Web site: “Why is it so easy to find out the server uptime, but so difficult to figure out the downtime?” And, it all seems so obvious in hindsight. According to the Web site:

downtimed is a program for monitoring operating system downtime, uptime, shutdowns and crashes and for keeping record of such events.

...*downtimes* is a command-line tool, which can be used to inspect previous downtime records recorded in the downtime database file.

Installation Source tarballs in both tar.gz and tar.xz format are available at the Web site. As far as library requirements go, I didn’t run into any hiccups during installation, and I couldn’t find anything about needed libraries, so the requirements seemed very minimalist. As for the actual installation, at least the first part is an easy process.

Grab the latest tarball from the Web site, extract it, and open a terminal in the folder. From there, it’s the usual:

```
$ ./configure
$ make
```

```
File Edit View Bookmarks Settings Help
nhoj@nhoj-desktop:~$ downtimes
down 2010-10-28 20:55:00 -> up 2010-10-28 20:55:29 = 00:00:29 (29 s)
crash 2010-10-28 22:55:54 -> up 2010-10-28 22:56:30 = 00:00:36 (36 s)
down 2010-11-01 15:50:24 -> up 2010-11-01 16:25:43 = 00:35:19 (2119 s)
nhoj@nhoj-desktop:~$
```

A program that should have existed years ago, *downtimed* is a very simple tool for recording computer downtimes.

If your distro uses sudo:

```
$ sudo make install
```

If your distro uses root:

```
$ su
# make install
```

From here, things aren’t so easy, and you should consult the documentation for further instructions, so I’ll leave the next steps to you at this stage. If you’re lucky, someone will have made a binary for your system, and in the case of Ubuntu, I found a custom package at the Ubuntu Personal Package Archives at launchpad.net.

Assuming you were able to finish installing *downtimed*, you’ll obviously need to reset to complete the process.

Usage Actually using *downtimed* is probably the easiest part, and it’s possibly the easiest program I’ve ever reviewed.

Open a terminal and enter:

```
$ downtimes
```

If everything went properly during the installation, you should have a record of your first shutdown, with a likely downtime of something under a minute on modern distributions. *downtimed* is clever enough to distinguish between a crash and a proper shutdown, which I tested (after saving all of my work) by pressing the computer’s reset switch. The following night we

had a power blackout of just a few seconds, and that also registered as a crash. If you look at the screenshot included, you can see that clean shutdowns are annotated to the left with a “down”, while others are marked with a “crash”.

That’s pretty much all there is to it really. Of course, the program can be extended in certain ways, such as specifying another database file instead of the default and viewing uptimes, but I’ll let you explore those for yourself in the program’s man pages.

In the end, this is one of those programs you can’t help feel should have existed already, but has come along only recently. I love the program’s simplicity, which in particular makes me think this will become a mainstream tool. Distro makers take note: this should be in your deployed systems by default, especially on server-style distributions.

Performous

performous.org

Describing this game is difficult. Under the heading of “What is it?”, the Web site gives this simple statement: “*Performous* is a free cross-platform music and rhythm game.” But, that doesn’t even begin to describe this project.

After several days, all I can come up with is that *Performous* is a cross between *Karaoke*, *Guitar Hero*, *Rock Band*, *Frets On Fire*, *SingStar*, *StepMania* and *Dance Dance Revolution*.

Here’s the feature list from the Web site, which should sum things up nicely:

- Should work on almost any platform.
- Very accurate singing pitch detection. Noisy environments are not an issue. Multiple simultaneous tones can be detected and separated properly.
- Intelligent gameplay with guitar and bass. The engine determines what was intended rather than simply picking the nearest chord. Hammer-ons and pull-offs are undone when the player accidentally does them and then picks the note in the regular way.
- Drums.

If You Use Linux, You Should Be Reading

LINUX JOURNAL



In one of the most bizarre experiences of my life, here's a DDR-style dance track to German metallers Rammstein! Only with OSS.



The open nature of *Performous* means alternative bands are as popular as the mainstream. Here I'm drumming to Finnish band Nightwish.

- Dance pads.
 - Device autodetection: *Guitar Hero* and *Rock Band* instruments, generic dance pads, *SingStar* and *Rock Band* microphones, and fallback to other audio devices when no SS/RB mics are detected.
 - Top-notch visuals: fully OpenGL-based rendering, music videos as backgrounds, smooth UI animations and other eye candy.
 - Solid C++ engine: performs well on slower machines too. Rapid but still safe development, thanks to advanced C++ features. Plenty of threaded stuff like background loading of songs and videos.
- Free software, licensed under GNU GPL version 2 or later.

Installation If you're chasing a binary package, there are all sorts available for Ubuntu (I recommend the cutting-edge packages from PPA at Launchpad), as well as Debian, openSUSE, Arch Linux and Mandriva. If you can get a package, do so. A source tarball is available, but it's a long process of chasing down dependencies. Nevertheless, in the interest of distro-neutrality, I'm running with the source as usual. But strap yourself in: we're in for a long one!

Luckily for me, the *Performous* documentation is quite thorough. As for library requirements, the docs say you need the following:



- » In-depth information providing a full 360-degree look at featured topics relating to Linux
- » Tools, tips and tricks you will use today as well as relevant information for the future
- » Advice and inspiration for getting the most out of your Linux system
- » Instructional how-tos will save you time and money

Get *Linux Journal* delivered to your door monthly for 1 year for only \$29.50! Plus, you will receive a free gift with your subscription.

SUBSCRIBE NOW AT:
WWW.LINUXJOURNAL.COM/SUBSCRIBE

Offer valid in US only. Newsstand price per issue is \$5.99 USD; Canada/Mexico annual price is \$39.50 USD; International annual price is \$69.50. Free gift valued at \$5.99. Prepaid in US funds. First issue will arrive in 4-6 weeks. Sign up for, renew, or manage your subscription on-line, www.linuxjournal.com/subscribe.

- OpenGL: graphics.
- Boost: threads, parsers/conversions, pointer containers.
- SDL: window management and input.
- FFmpeg (avcodec, avformat and swscale): A/V decoding.
- Cairo >= 1.2: scalable graphical support.
- librsvg: SVG rendering (themes).
- libxml++: XML parser, used for themes, etc.
- PortAudio v19: audio capture and playback.
- Gettext (optional): for internationalization.
- PortMidi (optional): for MIDI drum support.
- OpenCV (optional): for Webcam support.

However, I needed to install a lot of development packages, as well as some dependencies I didn't see listed before I could continue. As for my own system (Kubuntu), I had to install the following: libboost-all-dev, libglew1.5-dev, libavutil-dev, libavcodec-dev, libavformat-dev, libswscale-dev and libportaudio-dev.

Once you have the needed dependencies, download the latest tarball, extract it, and open a terminal in the folder. Now, enter the following commands:

```
$ mkdir build
$ cd build
$ cmake ../
$ make
```



From the promo shots: karaoke, guitar and drums all at once!

On *Performous*, I was able to play along to my favorite (and rarest) tracks, and from a drummer's perspective, the user-made tracks are actually a lot more realistic than the commercial ones.

If your distro uses sudo:

```
$ sudo make install
```

If your distro uses root:

```
$ su
# make install
```

Whether you're running from source or not, I *highly* recommend installing *Frets On Fire's* song files, as this gives the game a selection of songs to choose from immediately. I also found an Ultrastar song package at the PPA site, and if you can find similar packages, I recommend installing them.

To run the game, enter:

```
$ performous
```

Usage As the feature list suggests, the game *should* autodetect any of the devices from these popular rhythm/karaoke games. I don't have any of these devices, so I can't verify its efficacy, but perhaps a PlayStation-ist out there with all the equipment can tell me how it went.

If you're using the karaoke features, you'll want to test out your sound levels in the Practice section before moving on to the game itself. If you have only the basic *Frets On Fire* songs installed, for now you will have only guitar songs, so I'll cover those first.

The five frets you'll be using are controlled with the number keys 1–5, with either R-Shift or Enter as the picking button. However, given that you can use either button or both, I found that for super-fast sections it helps to thrash between both buttons with two fingers.

This may sound sacrilegious, but instead of picking up the keyboard, I found it easier to leave it in place on the desk, and just play it flat, like a piano. I know, I know, you're meant to pick it up to simulate a guitar or bass, but as someone who owns both a bass and a guitar, trust me, it doesn't feel anything like

either (not even with a GH controller).

However, you'll soon want to move on to other songs with other instruments (you won't even be able to do karaoke with that selection), so I recommend having a look around on the community boards for *Frets On Fire*, *SingStar* and *StepMania*. The *Performous* documentation says it looks for songs in several directories, but I had limited success. In the end, I just copied my new songs into `/usr/share/games/performous/songs`, and they worked without a hitch.

If you want to play a drum track, for instance, you need a song that's had a drum track programmed with it. Assuming the song files are all where they should be, it'll appear in your song menu the next time you start *Performous*. If a song has multiple tracks and you want to play only the drums, when the song is loaded, press one of the drum pads/buttons to enable the drum mode (or a dance pad button for dance mode and so on). As for bass, because it uses the same buttons as guitar mode, you need to change from the default guitar mode to bass mode at the start of the song.

On *Performous*, I was able to play along to my favorite (and rarest) tracks, and from a drummer's perspective, the user-made tracks are actually a lot more realistic than the commercial ones.

Ultimately, *Performous* is both an ambitious and a pragmatic concept. People are playing the same selection of games anyway, so why not combine them? And, given that *Performous* can use the real controllers, it means I also can skip getting a PlayStation (I'd rather eat my own shins than be a console player). Yay for PCs! ■

John Knight is a 26-year-old, drumming- and climbing-obsessed maniac from the world's most isolated city—Perth, Western Australia. He can usually be found either buried in an Audacity screen or thrashing a kick-drum beyond recognition.

Brewing something fresh, innovative or mind-bending? Send e-mail to newprojects@linuxjournal.com.

Attend the *FIRST*

Android Developer Conference

AnDevCon

The Android Developer Conference

March 7-9, 2011, San Francisco

San Mateo Marriott

AnDevCon is the technical conference for software developers building or selling Android apps

- Learn to build apps for Android, the fastest-growing mobile platform
- Learn to use the official Google tools, APIs and more
- Improve your skills for competitive advantage
- ISV entrepreneurs: Learn how to market your Android apps using its specific channels

Come take advantage of our **55+** technical classes and workshops!

- Create Effective Smartphone Apps for the Enterprise
- Android Internals: From SDK to NDK to APIs
- Creating and Using Secure and RESTful Enterprise Services
- Beyond English: Make Your Android App a Global Success
- Speed Up Android Dev with Model-Driven-Development
- Android for .NET/C# Developers Using MonoDroid
- Building Location-Based Services (LBS) Apps in Android
- Hello CouchDB: Building Android Apps on Top of CouchDB
- Building Rich Mobile Apps with HTML5, CSS3 and JavaScript
- Top Tips for Android UIs
- Your First Flash Application for Android
- Mobile Data Synchronization with Any Database
- Creating an Easy-To-Use Modular Framework
- Writing Reusable Android Components
- Securing Your Android Applications
- Test and Debug with Android Dev Tools
- Androidimation: Animation in Android
- Using HTML5 to Target Android
- Android 101
- The Six-Hour App

...and more!

Check out full class descriptions at www.AnDevCon.com

Follow us at twitter.com/andevcon

More Than
One Billion Apps
Downloaded!

Register Early
and SAVE!



HARDWARE

Giada Slim-N20

I've put everything from a full-tower PC to a homemade embedded mini-ITX computer in our living room connected to our television. They've all done their jobs, but nothing has looked as sleek as the Slim-N20. SHAWN POWERS

Giada's Slim-N20 is a Nettop computer designed with living-room entertainment in mind. Although it comes installed with a full operating system (mine came with Ubuntu 10.04 pre-installed), the lack of keyboard/mouse and inclusion of a remote control makes it clear this computer is meant for more than menial desktop chores. The Slim-N20 has its shortcomings, but it also has some awesome features hard to find elsewhere in the Nettop market.

Hardware

The specs for this little computer put it right in the middle of the pack when it comes to Nettop specs. My test model came with the following:

- Intel Atom D510 (two cores, 1.66GHz).
- NVIDIA GT218-ION (ION2, next-generation).
- 4-in-1 card reader (SD, MMC, MS, MS Pro).
- Three USB 2.0 ports.
- One eSATA port.
- HDMI + VGA.
- S/PDIF digital audio.
- Gigabit Ethernet.
- 320GB SATA2 hard drive.
- 2GB DDR2 RAM.
- 802.11n wireless networking.
- Bluetooth.

In addition to the internals, this little PC shipped with an infrared remote (Windows Media Center-compatible, more



Figure 1. Giada Slim-N20 and Components

on that later). It includes a solid metal stand, external power transformer, three-foot HDMI cable and DVD (although there's no DVD drive).

The connectors are arranged perfectly for a living-room PC, but if the computer is going to be used as a desktop workstation, a USB port or two on the front would have been nice. Figure 2 shows the connectors on the back of the unit: power, Ethernet, USB, S/PDIF and HDMI. The top of the N20 has a flip-open area concealing the analog audio in/out, combination eSATA/USB port and 4-in-1 card reader (Figure 3).

It's frivolous, but I think my favorite part about the hardware is the power button. It's turned on by squeezing the unit in the front, and the LED indicators are multifunctional. You can see hard drive activity by the coloration and blinking of the crescent shaped ring around the power button. Yes, I realize it's a silly

thing, but I found it rather cool (Figure 4).

Performance

The Giada Slim-N20 is touted as a media playback device and gaming unit. The Atom D510 CPU really makes this unit a tough sell for gaming, but with the ION2 graphics chipset, it performs moderately well on simple games. I'd consider it a Nettop computer that I could waste a few minutes gaming with, but I certainly wouldn't consider it a gaming machine.

The ION2 GPU does really do its job as an HD-decoding beast, however. My tests with 1080p video played without a stutter. I wasn't able to stream 1080p video very well over the 802.11n network, but I really wouldn't expect to be able to stream full HD over wireless anyway. The gigabit Ethernet port was able to keep up just fine with any network-shared video I threw at it, whether they were MKV files



Figure 2. The back of the computer has most of its connections.



Figure 3. The analog audio, eSATA and card reader are a little difficult to get at, but if not needed, they're well hidden.




Figure 4. Maybe it's just the kid in me, but I really like this power/activity light.

It's frivolous, but I think my favorite part about the hardware is the power button.

The menus are responsive, programs start quickly, and most important, every bit of hardware worked out of the box. In comparison to the Acer Veriton I've been running as a second Linux desktop, the Giada Slim-N20 was much nicer.

User Experience

I won't go into very much detail regarding the N20's user experience as a desktop machine. I do wish it had a few more USB ports, but apart from that, it's a surprisingly zippy little machine, especially considering it's a Nettop. Where this unit really shines, however, is in the living room.

With all major connectors in the back, the N20 hides the fact that it's a computer very well. With normal operation (and digital audio), the top panel never needs to be opened. The only thing visible from the front is the very thin (less than 1") profile and the cool LED power/activity button.

The N20 is also quiet—*very* quiet. Even when churning out 1080p video, the unit never was audible, even during quiet scenes in the movie. The trade-off for such quiet operation is that it does get a little warm when under a load, but even that isn't worrisome. It gets warm to the touch, but never hot.

The included remote works. It's not great, and it's basically a Windows Media Center remote, but because the MCE remotes are so popular, XBMC and other programs can deal with its awkward layout fine. The volume buttons also work with the Ubuntu desktop, which makes it possible to turn the volume up and down with a single button push, which is nice if your keyboard doesn't have multimedia keys.

One last note about the desktop experience—the unit I was shipped came with Ubuntu 10.04 installed, and it was stock. There were no strange interfaces, awkward menu systems or branded splash screens. It was just plain Ubuntu, which suits most users just fine. Giada, thank you for not forcing your logo down our throats at every boot.

or H.264 MOV files. I've read others had problems with 1080p video on the N20, but perhaps because I was using XBMC Live, which is tweaked for the NVIDIA GPU, I had better luck.

I didn't do any hard-core benchmarking, because quite honestly, those numbers are easily accessible on the Web. I figured for this review, it would be more helpful to compare the N20 to other Nettop machines I've used. In short, it was great. Startup time from the internal 320GB SATA drive was impressive. I was able to use the system as a desktop machine without noticing it was a Nettop.



Figure 5. On the left is a standard Windows MCE remote and on the right is the remote that comes with the N20. They're very similar.

The Good

The only thing about this Nettop that excited me before doing this review was that it had the ION2 video chipset. I was fairly certain it would allow programs like XBMC to play just about any video I threw at it, and my suspicions were happily confirmed. What I wasn't expecting was how nice this little computer *looks*. It's amazingly thin and designed with subtlety in mind. Generally, Apple does the best job of making a computer look like a fashion

accessory, but I must admit, the Slim-N20 looks great next to my television. install (well, I used XBMC Live for the living room), but I would expect no problems with other distros.

The Bad

The limited number of USB ports is a bit frustrating, considering a keyboard and mouse take up the only two exposed ports in the back. Flipping open the top panel to access the combo USB/eSATA port isn't the end of the world, but at least one additional port on the front would be welcomed. My unit also didn't

The N20 is also quiet—very quiet.

When you add the impressive desktop performance, the N20 really did surprise me. It's not a workhorse by any means, but for a unit smaller than most thin clients, it makes for a quite tolerable user experience. If you do a lot of video rendering or hard-core compiling, I'm sure the small CPU will frustrate you. For many users who would consider a Nettop in the first place, however, the N20 is fun. Because the components are very Linux-friendly, installing a replacement distribution shouldn't be a problem either. I reviewed only the included Ubuntu

auto-detect that I was trying to use HDMI instead of VGA. That's not a showstopper, but it was annoying when I tried to hook it to my television.

The included remote, although nice, obviously was designed for Windows Media Center. I know the Slim-N20 also ships with Windows, but the MCE remotes are so awful, I wish they would have an alternative for those of us not interested in Windows (Figure 5). The remote also feels cheap, in stark contrast to the very sturdy feeling of the computer itself. Buying an additional remote is a possibility, but there's always the chance it won't work with

the infrared receiver and so on. Plus, buying an additional remote would add to the price, which brings me to my next point.

The Ugly

The Slim-N20 is expensive. There's really no way to claim it's not. It's arguable that the features and design make it worth the \$450, but only if you really value aesthetics. You certainly could get a more powerful desktop machine for the money, but it wouldn't be nearly as nice sitting in your living room. If things like power usage, case design, quiet performance and beauty are important to your purchase, the Slim-N20 won't let you down. If you're just looking for a Nettop device to browse the Web and e-mail your grandparents, you might be happier with a cheaper model.

The Verdict

I'm actually impressed with the Giada Slim-N20 Nettop computer. It's not a cheap, throwaway device by any means, but if you're looking for a stylish media-center computer that is fully customizable, it's hard to beat. If your living room and television double as your workstation and office, it might get you the best of both worlds.

It's not the fastest computer for the price, but it's my favorite Nettop unit to date. If you don't mind spending a little more money for a fashionable computer that doesn't come with pre-installed garbage and bloatware, check out the Slim-N20. I'll be reviewing the Boxee Box by D-Link in an upcoming issue. Although it doesn't act as a desktop computer at all, it's about half the price of the N20, and it's designed for the living room as well. ■

Shawn Powers is the Associate Editor for *Linux Journal*. He's also the Gadget Guy for *LinuxJournal.com*, and he has an interesting collection of vintage Garfield coffee mugs. Don't let his silly hairdo fool you, he's a pretty ordinary guy and can be reached via e-mail at shawn@linuxjournal.com. Or, swing by the #linuxjournal IRC channel on Freenode.net.

Resources

Giada: www.giadatech.com

My Video Review of the Giada Slim-N20: www.linuxjournal.com/video/giada-slim-n20

Jolicloud

Legacy hardware, man up with Jolicloud. MIKE DIEHL

I once had a frustrated customer point to his computer and exclaim, “I don’t want to know how it works. I just want to use it!” I know what he wanted; he wanted to push a button and have the machine do what he expected. Over the years since that episode, computers have become a bit more friendly and easy to use, but they’ve also grown in footprint and complexity, or at least, that’s what I thought. Finally, I think I’ve found an operating system that is powerful enough for everyday use, yet easy enough to use that almost anyone can learn to use it. Jolicloud is based on Ubuntu and is targeted toward Netbooks and other low-power devices. I was able to base this review on a pre-release image of version 1.1, which should be available by the time you read this article. The release notes indicate that version 1.1 includes an upgrade to the underlying Ubuntu from 9.04 to 10.04 as well as several package upgrades and kernel optimizations. The release notes also claim improved battery life, system stability, supported hardware and overall quality of the operating system. Additionally, I’ve been told that several changes and improvements already have been implemented since I downloaded my snapshot for evaluation, so your mileage may vary, although only slightly. I was pleased to discover that, in typical Ubuntu fashion, almost everything seemed to work right out of the box.

For this review, I tested Jolicloud using a VirtualBox virtual machine with 512MB of RAM. I started with a live .iso image and was able to become familiar with Jolicloud before actually installing it. The .iso image is the most convenient way to get Jolicloud onto a workstation, provided it has a DVD player. However, using the USB key installation method probably is more practical for most Netbook owners. The Windows installer allows Jolicloud to be installed side by side on devices that have Windows installed. So, simply testing Jolicloud doesn’t really require any type of commitment. During the testing phase, Jolicloud gives you the option of installing onto your system’s hard drive. The entire installation process takes about 30 minutes, and nothing could be easier.

Once the installation was complete, I had to create an account at my.jolicloud.com, so I

then could associate my new machine with that account. The idea is that you can associate more than one machine with a single account, and they’ll all be synchronized (more on this later). Once this association was made, I was able to log in using the user name and password I had set up during the installation process. Alternatively, I could have tied my user account to my Facebook account and used those credentials to log in.

Once users are logged in, they are presented with a clean, simple interface. Figure 1 shows Jolicloud’s Launcher. As you can see, I’ve gone ahead and installed a number of applications. The application icons are large and obviously geared toward touchscreen users. The Launcher presents the application icons one page at a time, and each page can be selected by clicking one of the little dots at the bottom. Yes, it’s very much like using an iPhone.

In the upper right-hand corner is a simple control panel, just like you would expect on most systems. In the upper left-hand corner, you see a small blue cloud icon and a couple other icons. This is the taskbar. Any task can be brought to and from the foreground by clicking on its icon. The blue cloud always will be there. Clicking it minimizes all windows and shows the home screen. Once you’re at the home screen, you will see four buttons at the top. The first one on the left brings up the Launcher. The button that sort of looks like a satellite dish allows you to follow other Jolicloud users and see what applications they are installing and which ones they like. The button that looks like a file folder brings up a simple file manager. From here, you can manage files on your local storage, as well as files stored in various



Figure 1. Jolicloud’s Launcher

cloud providers, such as Dropbox.

Finally, the button that looks like a gear allows you to perform miscellaneous functions regarding your Jolicloud account and your registered machines. This is also where you’ll find the Local Settings application.

Although Jolicloud isn’t extensively configurable, some gems can be found under Local Settings. Under Network Connections, you can configure various wired, Wi-Fi, 3G, DSL and VPN connections. The VPN capability includes Cisco, PPTP and OpenVPN clients. Local Settings also is where you’ll find network tools, such as ping, traceroute and nslookup. Local users can be added, deleted and modified as well.

Going back to the home screen, notice the big green Add button in the upper-left corner. This button takes you to Jolicloud’s application store (Figure 2). Along the left-hand side, you will see that you can search for applications by category as well as by “Most Liked” and “Most Recent”. I don’t know how many applications are available for Jolicloud, but I’ve spent hours browsing the collection. When I first saw that Facebook was being offered as an “application”, I thought that it was sort of cheating. Of course, I was under the mistaken impression that the result was simply a browser bookmark. It turns out that the resulting application doesn’t actually launch a full-up browser. Instead, it launches a Chromium runtime environment and the Web site ends up looking like a standalone application. Because of this capability, Web-based (or cloud-based) applications look just like standard applications. As long as your machine is connected to the Internet, there is no difference between running a Web-based application and a standalone application. For example, the Wikipedia cloud-based application behaves just like the locally installed Celestia application.

This blurred distinction between cloud and local applications is a two-edged sword. On the one hand, it offers a vast amount of functionality with little or no software installation. On the other hand, more naïve users could become frustrated when their cloud applications don’t work when their computer is off-line. Jolicloud takes full advantage of cloud-based computing by blurring the distinction almost perfectly. But, with the addition of cloud-based storage like Dropbox, this blurring extends beyond applications to file storage. Jolicloud makes such cloud-based storage spaces look just like local storage. While I was evaluating Jolicloud, I went ahead and installed the



Figure 2. Jolicloud's Application Store

Dropbox application and applied for a Dropbox account. Once this was configured, I had a Dropbox directory in my home directory and Jolicloud's file manager treated it just like any other directory.

The search field at the top of the home screen is pretty novel. Once you enter a search term, you are presented with a pick list that allows you to use the search term to search Google, Yahoo and Bing, as well as Facebook and Twitter. You're also offered the opportunity to use your search term to search for Jolicloud applications and users.

Finally, there's the feature I alluded to earlier, synchronization. In this case, synchronization refers to the fact that if an application is installed on one machine, it automatically will be installed onto each machine associated with the my.jolicloud.com account. This way, if you had Jolicloud installed on a laptop, a workstation, a VM and a tablet, you could install an application on one device and find it installed automatically on all of the rest.

During my time playing with Jolicloud, I was able to see this feature firsthand. After I had been using Jolicloud for some time, I decided to create another virtual machine instance. This installation went as well as the first ones. When I logged in to the new machine, I saw icons for all the applications that had been installed on the other machine, only slightly dimmed. In a few seconds, the new machine started to download the corresponding applications. From a nerd's point of view, it was pretty neat to watch, and it was completely automatic. Of course, none of these features would be of any use without applications.

Jolicloud's application store allows you to install many applications with just a few mouse clicks. The Chromium Web browser is installed by default, but you can install Firefox if that is your preference. For serious office work, Jolicloud offers Google Docs as well as the OpenOffice.org suite. Various communications tools also are available, including

Pidgin and Skype. Even fun tools like Boxee and Celestia are easily installed.

But, Jolicloud isn't for everyone. Power users and developers may find Jolicloud a bit limiting—especially given the minimal number of supported screen resolutions and lack of availability of developer's tools in the application store. Even with these limitations, however, Jolicloud could support a broad base of users.

People who don't want to carry around a bulky laptop but who still want the speed and utility of a fully loaded Linux system could install Jolicloud on a Netbook and have just about all of the functions they could want. Once installed, such a system certainly would out-perform similar hardware running the factory-installed Windows operating system. I'm sure that the ability to stream media, update Facebook, use an instant-messaging client, and do homework using a complete office suite, without bogging down the computer, would be welcome by any student.

Referring back to the customer I mentioned at the beginning of this article, many people just want to turn a machine on and be able to do useful work with it. Those people don't want to perform any incantations to load a document, and they don't want to have to configure anything. Jolicloud's simple user interface makes it ideal for people who don't want to be computer geniuses and just want to use a computer.

People who appreciate simple interfaces and large icons also would like Jolicloud. Additionally, those people may not be as patient with slow system response, but at the same time, they may not be able to afford modern equipment.

Finally, I could see deploying Jolicloud in a student computer lab. Assuming that all of the computers were associated with the same account, maintenance would be quite easy. Software that was installed on one machine would propagate to all of the other computers, and virus protection wouldn't be needed as it is with Windows-based systems. Using Linux's inherent permissions system, these machines could be secured against tampering easily. Once again, Jolicloud's meager resource requirements mean that schools could continue to use older equipment.

Overall, I've been very impressed with Jolicloud, but it can be improved. A more-sophisticated permissions system could be devised. Such a system could be used to prevent student users from installing applications or changing the network configuration. This would make Jolicloud machines much more

secure in a laboratory environment. A remote user-authentication mechanism, outside of Facebook, also would make the management of large Jolicloud deployments much more convenient. With so many applications available from Jolicloud's application store, I still wish I had some way to install my own applications on the device and have corresponding icons appear on the home screen. Granted, I might be able to find where the home screen is stored and hack it in order to add icons, but there must be a better way—or there should be.

With as many communications applications as Jolicloud has available, I was surprised there wasn't a VoIP client other than Skype. I also was surprised not to find an SSH server running.

Finally, because Jolicloud is so ideal for portable devices, it would be nice if there were some facility to track or disable the device remotely. Perhaps devices that had an integrated GPS could be configured to post their location periodically. Or they could use a dynamic DNS client to make their current IP addresses known. It's probably just a matter of time before most of these weaknesses have been resolved, as most of them simply require the appropriate software to be written or made available to the Jolicloud community.

Although I'm not quite ready to toss out my KDE environment in favor of Jolicloud, I do expect to be using Jolicloud for a few things. I've got an old laptop that never was really powerful enough to run Windows XP, but it certainly should be able to run Jolicloud. I'm also thinking seriously about installing Jolicloud on my young sons' computer in a dual-boot configuration with the existing Vista. Most of the school-related work that they need to do can be done within Jolicloud, and that might reduce some of the temptation to play games when they should be doing their schoolwork. Being able to boot in to Vista to play graphics-intensive games then could be used as a reward for finishing homework on time.

I hope this article has enticed you to try Jolicloud, either on a Netbook, desktop or even in a virtual machine. If you do, I think you'll find that Jolicloud has a very elegant interface, a wide variety of applications and is a great way to bring Linux's stability and performance to otherwise-limited computing devices. ■

Mike Diehl is a contract programmer and consultant in Albuquerque, New Mexico. Mike lives with his wife and three small boys and can be reached via e-mail at mdiehl@diehlnet.com.

Ubuntu 10.10 Maverick Meerkat in Amazon EC2

Don't be afraid of the cloud! Let Canonical and Amazon pick up the tab. BILL CHILDERS

Last year, I wrote an article on how to run Ubuntu 9.10 in the Amazon Elastic Cloud (LJ, May 2010). But, the folks over at Canonical haven't been sitting on their laurels during the past year. They're so proud of the work they've done on their latest release, 10.10 Maverick Meerkat, they're offering a free 55-minute trial of it in Amazon's EC2 cloud. The best part is you don't need an Amazon EC2 account or a credit card. The wizards at Canonical have tied this into their Launchpad software collaboration platform, so all you need is an account on launchpad.net to "get your cloud on".

What you get for your 55 minutes is actually a fairly nice demo package. Your cloud instance consists of the equivalent of a 1.2GHz processor, 2GB of RAM, 160GB of disk space, full sudo-enabled root access and a fast local connection to the Ubuntu mirrors. The Ubuntu folks even provide a few pre-configured applications from which to choose, like WordPress, MoinMoin or Drupal 7. The really great part about this setup is that it renders most of my previous article obsolete—it's really easy to set up and launch your instance.

Getting Started in the Cloud

Before you log in to the "Ubuntu in the Cloud" Web site, you should generate a set of SSH keys on the client (or set of clients) you'll use to administer the cloud instance, and upload those to your account on launchpad.net. This is a prerequisite, because Amazon EC2 generally allows SSH access only via SSH keys. If you already have a launchpad.net account, it's easy to do. Just go to your account page, click on SSH Keys, then click on the green plus icon to add a new SSH key. Then, paste your key in the field. If you've never generated any SSH keys before, there's a great tutorial on the Ubuntu Wiki (see Resources). If you don't have an account on launchpad.net, you should.

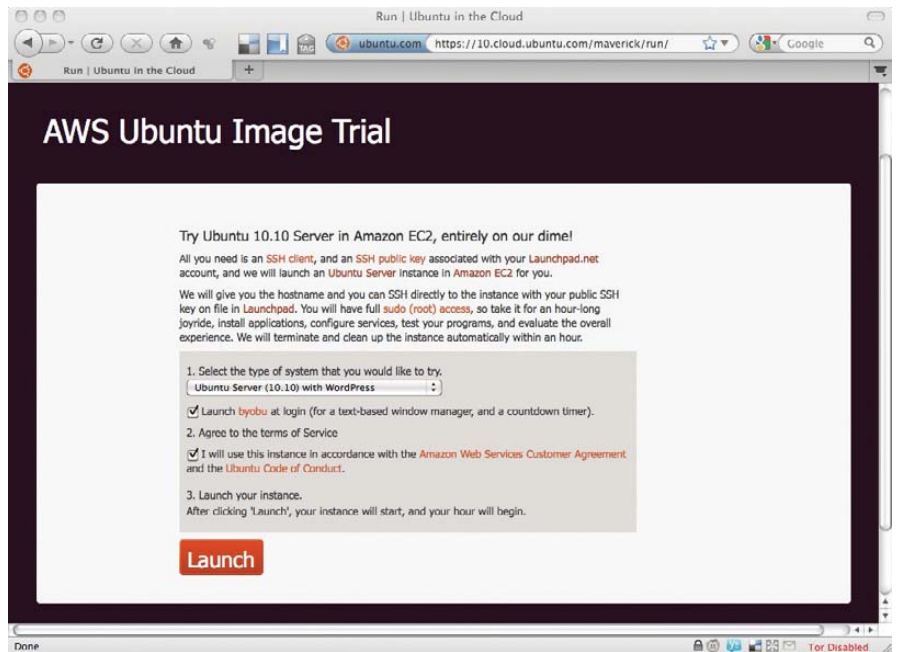


Figure 1. Launching Your EC2 Demo Instance

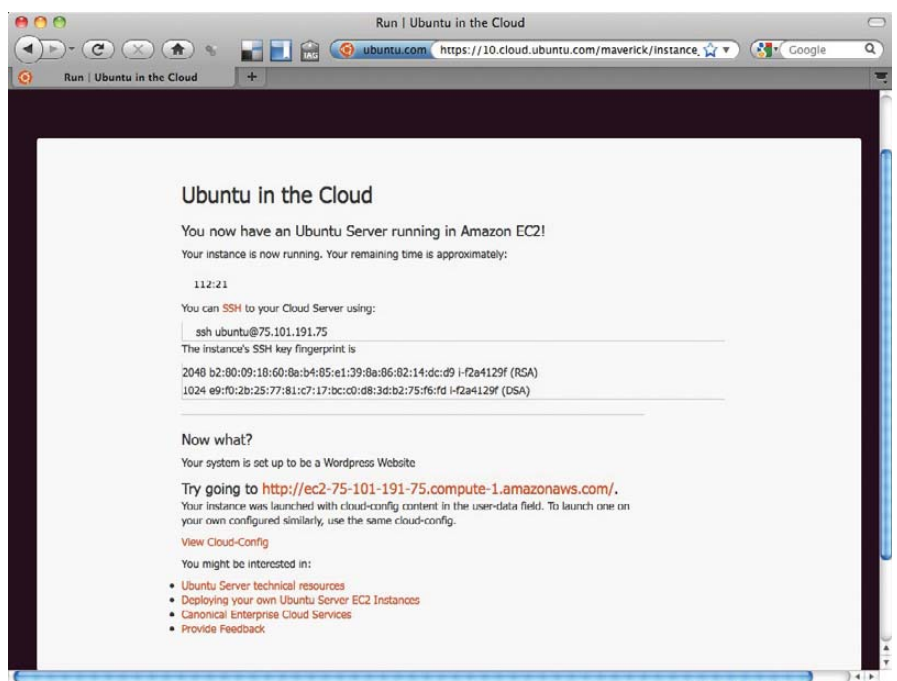


Figure 2. How to Access Your EC2 Demo Instance




Figure 3. SSH'd to my new instance, with /proc/cpuinfo displayed. Note byobu at the top of the window.

It's where you can file bugs and enhancements to Ubuntu and other projects, and it's easy to register for one.

Once you've uploaded an SSH key to launchpad.net, point your Web browser at the "Ubuntu in the Cloud" link (see Resources), and click the Try Ubuntu 10.10 button. You'll be prompted to enter your launchpad.net credentials if you're not already signed in to Launchpad. Once your credentials are accepted, you'll see the options for your

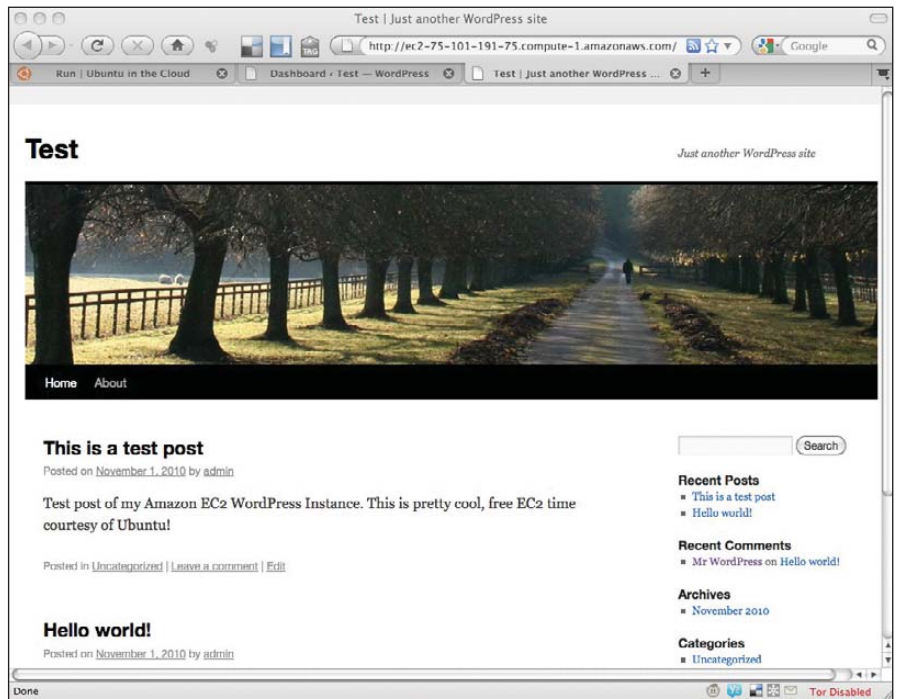


Figure 4. WordPress is live! Easiest installation I've ever done.

Now that your instance is on its feet, you can do anything you'd do with an Internet-connected Ubuntu server—so long as you are done within an hour.

free EC2 instance. Select what you'd like to try (Base install, WordPress, MoinMoin or Drupal) and whether to use byobu (a screen-based terminal wrapper—I highly recommend it). After that, just agree to the Terms of Service, and click the big orange Launch button to start your instance. You'll see an Ubuntu logo with a little status indicator blinking away while the instance is built.

When the instance is ready for use, you'll see a countdown clock in your browser with the remaining time and the IP address of your instance. You can SSH in to your instance by logging in as "ubuntu@<ip address>", or if you selected one of the preloaded software packages, like WordPress, the Web page will display a link you can use to get to the admin interface of your package.

Now that your instance is on its feet, you can do anything you'd do with an Internet-connected Ubuntu server—so long as you are done within an hour. You can configure it as a Web server, play around with LDAP, make it a mail server, or even fire up a honeypot (so long as you're done with it in an hour). Anything you can apt-get in

Ubuntu, you can install in this server, so play with things like Apache, Squid or whatever servers you like. This is a full Ubuntu installation in Amazon's data center, so be sure to test-drive it!

The preconfigured software is pretty good as well. In this case, when I provisioned the instance, I selected the WordPress option, and I was pleasantly surprised to see WordPress ready to go. All I had to do was click the URL presented by the Web page once the instance was ready and answer a few WordPress-specific questions like "What do you want the site name to be?" and "What's the admin password for your

WordPress site?" I literally had posted my first blog post to the instance within a couple minutes of it being live.

As your hour winds down, if you're SSH'd to the instance, you'll get a "wall" message ten minutes before your hour is up, another when there are five minutes remaining, and a final message at the one-minute mark. Once your time is up, the instance "powers off", and the Amazon cloud automatically cleans up and deletes your instance, as well as any data you put there, so make sure you're not putting anything important there. It truly is a demo service.



Figure 5. Oh no! I'm melting—my EC2 instance, shutting down and deleting itself.

Okay, I like this. But it's gone in an hour, and I want more.

So, you've played with your server, installed other software, and monkeyed around with the preconfigured stuff on your server. You're hooked and want your own server, but you don't want to pay the full instance price of approximately \$70 a month to run a server full-time. Not to worry! Amazon has you covered. It now offers a completely free "micro" server, called the AWS Free

Table 1. Free Test-Drive Instance vs. Free Micro Instance

	Free Test Drive Instance	Free Micro Instance
CPU	1 EC2 Compute Unit (1.2GHz Xeon)	Burstable to 2 EC2 Compute Units
Disk	160GB of local instance storage	10GB of Amazon elastic block storage
Memory	2GB RAM	613MB RAM

Usage Tier, just for new users. The micro instance isn't quite as powerful as the free test-drive server, but it's absolutely serviceable for a blog, mail server or other light-duty application. Another possible (and very useful) application would be an off-site Nagios or other monitoring instance. See Table 1 for a comparison of the free test-drive instance vs. the micro server.

Unlike the free test-drive server (which is only 32-bit), the micro instance can be either 32-bit or 64-bit, though the low RAM of the system

and 1GB of storage.

- 100,000 requests of Amazon Simple Queue Service.
- 100,000 requests, 100,000 HTTP notifications and 1,000 e-mail notifications for Amazon Simple Notification Service.

The service does require a credit card to get started, but you're billed only if you exceed the usage outlined above. It's a great way to start learning

way to dip your toe in the water at no cost to you. Likewise, if you're familiar with the cloud, but you've not run Ubuntu before, give this test-drive a shot. You'll get enough time with Ubuntu to decide whether you want to pursue it further, but you won't have to take the time to spin up a machine of your own. Just let Canonical pick up the tab for the demo.

Once you've made up your mind as to whether the cloud is for you, take a look at the new AWS Free Usage Tier. For a personal server, sandbox or off-site monitor,

The micro instance isn't quite as powerful as the free test-drive server, but it's absolutely serviceable for a blog, mail server or other light-duty application.

doesn't really make that distinction very useful. The Free Usage Tier includes some other useful services as well, measured on a monthly basis:

- 750 hours of EC2 running Linux/UNIX on a micro instance (this is the server mentioned above).
- 750 hours of elastic load balancing plus 15GB data processing (in other words, you can load-balance between servers, if you've spun up another instance, but you'd pay for that second instance).
- 10GB of elastic block storage (includes one million IOs, 1GB of snapshot storage, 10,000 snapshot get requests and 1,000 snapshot put requests).
- 15GB of inward bandwidth and 15GB of outward bandwidth aggregated across all AWS services.
- 5GB of Amazon S3 storage, 20,000 get requests and 2,000 put requests.
- 25 Amazon SimpleDB Machine hours

how cloud services work and what they can do for you.

Conclusion

If you've never tried any cloud services before, or if you've been leery of giving your credit card to a provider, try the free Ubuntu 10.10 server. It's a great

the micro instance is up to the challenge, and the price definitely can't be beat. ■

Bill Childers is an IT Manager in Silicon Valley, where he lives with his wife and two children. He enjoys Linux far too much, and probably should get more sun from time to time. In his spare time, he does work with the Gilroy Garlic Festival, but he does not smell like garlic.

Resources

Ubuntu in the Cloud (Official Ubuntu/EC2 demo site): <https://10.cloud.ubuntu.com>

Try Ubuntu 10.10 Server in the Cloud for Free (blog announcement): ubuntu-smoser.blogspot.com/2010/10/try-out-ubuntu-server-1010-on-ec2-for.html

Launchpad Software Collaboration Platform: launchpad.net

How to Generate SSH Keys under Ubuntu: <https://help.ubuntu.com/community/SSH/OpenSSH/Keys>

Running Ubuntu 9.10 under Amazon's Elastic Cloud: www.linuxjournal.com/magazine/running-ubuntu-910-under-amazons-elastic-cloud

Amazon EC2 Pricing: aws.amazon.com/ec2/#pricing

Amazon EC2 Free Usage Tier: aws.amazon.com/free

COMING SOON to Linux Desktops

What's in store when the GNOME Project finally releases GNOME 3, and what will Canonical choose for the new Ubuntu?

Charles Olsen

Linux is great for people who like having choices. For example, if you want to use a graphical desktop, there are many from which to choose, including XFCE, LXDE, Fluxbox, Openbox and Enlightenment, not to mention the two most popular desktops, GNOME and KDE. See Shawn Powers' article for more information on other desktop options (page 48).

KDE version 4.0 was released in January 2008, but GNOME hasn't had a major version update since 2.0 was released in 2002. Quite a few incremental updates have been made to GNOME, and until recently, the GNOME Project had planned to release version 3.0 in September 2010. At the time of this writing, the release has been pushed to March 2011.

When KDE 4.0 was released, it was a major change from KDE 3.5. Although KDE 4.0 was labeled a stable release, it actually was intended for developers and early adopters. Much of the functionality from KDE 3.5 had not yet been re-implemented. KDE 4.0 felt like a huge step backward, and it was widely criticized as being incomplete. GNOME 3's delayed release may be an attempt to avoid a similar situation.

When I started using Linux, I came from the world of

Windows and preferred the KDE desktop because it seemed more familiar. I initially was excited to see the new KDE 4.0. But after the upgrade, I soon became disgusted with the mess I now had on my system. So I installed GNOME and was quite happy with it. Not only did it *work*, but also my PC ran faster with GNOME than it did with KDE 4.0.

I haven't completely ignored KDE 4 since then. I've occasionally installed it on a test system and tried it out. KDE 4.1 seemed generally usable, and by the time 4.4 was released, it actually seemed pretty good. But for my production systems, I've stuck with GNOME and have been quite happy with it.

Now, GNOME is close to releasing a major upgrade to GNOME 3, which uses GNOME Shell to display the application windows and other objects. Will this be another fiasco like KDE 4? I don't think so. Based on what I've seen so far, I think GNOME 3 will be a solid release.

But, that's not all that's coming. Canonical is charting its own path and developing a different desktop called Unity for its next version of Ubuntu. Unity is already Canonical's desktop for Netbooks—is it suitable for laptops and PCs?

GNOME 3

Although it has not been released officially, GNOME Shell has been in the repositories for a while now. You can install it and get a preview of what's coming in GNOME 3.

According to the GNOME Project Web site, "Our goal in designing GNOME Shell is to provide a consistent, self-teaching user interface based around the day-to-day tasks of the user." Let's take a look and see how GNOME plans to accomplish that.

Instead of the two panels in GNOME 2, GNOME Shell has a single panel at the top of the screen. It contains some of the items you're familiar with, such as the clock and notification icons. But instead of the three menus, now there's a single button named Activities. Next to the Activities button is the name of the currently running application. This eventually will contain the application menu and options, such as close, open a new window and options specific to that application.

When you click the Activities button, the full screen switches to the Activities overview mode. You also can invoke the Activities overview by pressing the Super key on the keyboard, or Alt-F1, or just by moving the mouse pointer to the upper-left corner of the screen. I hope GNOME offers an option to disable that last choice, because I found it way too easy to activate the overview unintentionally while working on my laptop.

The Activities overview takes the full

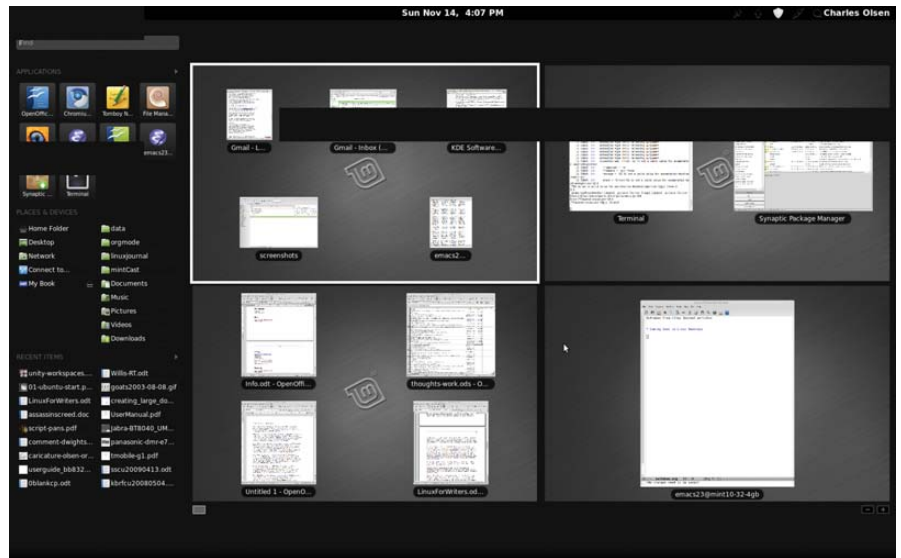


Figure 1. The GNOME Shell Activities Overview with Four Workspaces

a lot of applications open as I switch from one task to the next, and it's nice to be able to group the windows into a workspace so I can look only at the programs I'm currently using.

GNOME Shell lets you create workspaces as needed. When you're in the Activities overview, you'll see a small plus sign at the lower right. Click the plus sign to create a new, empty workspace. While still in the overview, you then can drag windows from the other workspaces into the new workspace. Drag an application icon into a workspace to launch the program in that workspace. You also can create

Pressing Alt-Tab displays all applications that are open, even if they're not in the currently selected workspace. A vertical separator bar separates the programs in the current workspace from those in other workspaces. You can select any program in the traditional way, by pressing the Tab key and releasing Alt when you get to the desired application, or you can use the mouse to click the application you want.

A new function that is not implemented yet is that Alt-` (the key above the Tab key on US-layout keyboards) will switch between windows of the same application. Some applications do

Canonical is charting its own path and developing a different desktop called Unity for its next version of Ubuntu.

screen, and it shows all the ways that users can switch from doing one activity to another. It displays a preview of all the windows that are open, on all the workspaces in use. It also shows a list of favorite applications, applications currently running, places (favorite directories and connected devices), and the 20 most recently used documents. There's also a Search box at the top left; if what you're looking for isn't immediately visible, simply start typing part of the application name or description.

I love workspaces. I tend to leave

a new workspace and launch a program in that workspace by dragging the application icon onto the new workspace plus sign icon.

In GNOME 2, switching to a running application in another workspace often meant switching to the appropriate workspace, then selecting the program you want. In GNOME Shell, the Activities overview shows all workspaces and previews of all windows in each workspace. I can go directly to the workspace and application simply by clicking the window in the overview.

this now with Ctrl-F6, but that takes two hands and is not consistent across all applications. Using Alt-` will be much easier and feel more natural.

As you use Alt-Tab to cycle through open programs, programs with multiple windows open will be grouped in one icon. If you pause while that icon is highlighted, a preview of the open windows will appear below the icon. You also can get the preview immediately by pressing the down-arrow key. You then can select the specific window by using the left- and right-arrow keys.

Customizing your favorite applications is easy. Use the Search box to find the application you want to add, then simply right-click the icon and choose Add to Favorites. To remove favorites you don't need, right-click the icon and choose Remove from Favorites.

The Places section of the Activities overview shows the items in the Places section in Nautilus. Simply add or remove items in Nautilus to customize the Places in your Activities overview.

The plan is that GNOME Shell will put a bigger emphasis on applications than the separate windows of the application. When you switch to the Activities overview, applications with multiple windows open will be represented by a single icon. In GNOME 2, you see an icon in the panel for each open window, and sometimes you have to guess which is the window you want.

In GNOME Shell, you can right-click on the application icon to see a list of open windows. You can select the specific window you want or open a new window. This wasn't working properly in the GNOME Shell I was running, so I wasn't able to test it.

The Search box will search not only application names and descriptions, but also through the names of recent documents and system customization commands.

When you're in the Activities overview with the small window



Figure 2. The Unity Desktop with Firefox Maximized

Ctrl-Alt-Right or Ctrl-Alt-Left, and you can move the current window to another workspace by pressing Ctrl-Alt-Shift-Right and Ctrl-Alt-Shift-Left.

GNOME Shell offers a new way to look at your desktop and activities, without making a radical break from the previous version on GNOME.

Unity

It would seem ideal for Ubuntu's April 2011 release to include GNOME 3.

However, Mark Shuttleworth has

launchers on the panel for quick access to their most-used programs. The new interface would follow that example and make a few applications instantly accessible, while still making it easy to get to everything else.

They also focused on getting the best use out of screen real estate. Unity initially was designed for Netbooks, where the screens were usually wide but not very tall. Now laptop and desktop monitors are moving to more of a wide-screen format. They realized they

According to Shuttleworth, the Canonical folks spent quite a bit of time analyzing screenshots of a couple hundred desktop configurations from the current Ubuntu and Kubuntu user base to see what people used most.

previews, you can zoom the preview of a particular window without selecting it by pointing to it with the mouse and rolling the mouse scroll wheel upward. Scroll down with the wheel to zoom back out.

Several of the functions still work the way you're used to from GNOME 2. Alt-F2 still brings up a command line where you can enter a single command. The Esc key will escape from anything, such as the Activities overview, the Search within the overview, and the Alt-F2 command line. You still can switch to a different workspace using

announced that Canonical will be taking a different direction and the next release of Ubuntu will feature Unity as the desktop.

According to Shuttleworth, the Canonical folks spent quite a bit of time analyzing screenshots of a couple hundred desktop configurations from the current Ubuntu and Kubuntu user base to see what people used most. They also wanted to identify things that were *not* needed in a lightweight environment.

They found that most users have between three and ten application

needed to be very conservative in using the vertical space.

Finally, they wanted the interface to be finger-friendly. Touchscreens are becoming more common, and Canonical wants Ubuntu to be ready.

This research and these decisions were the basis of Unity in Ubuntu Netbook Edition 10.10. And, according to Shuttleworth's announcement, Unity will be the default desktop for future editions of Ubuntu.

Rather than the traditional panels at the top and bottom of the screen, Unity

conserves vertical real estate by moving the bottom panel to the left side of the screen. This panel also is widened to make it more touch-friendly. It will show a few icons for instant access to selected applications, and it also will display icons for all applications currently running. Programs currently running are shown with a small indicator on the left side of the icon. The application that has the focus has an additional indicator on the right side. A single touch or click is all that is needed to launch a favorite application or switch to another program that's already running.

The three-menu design of traditional GNOME has been replaced with a single global menu that you invoke by clicking a button on the left panel. The titlebar, including the Close/Minimize/Maximize buttons of the currently selected application, is rendered in the top panel.

When you open the global menu, the top panel displays a search field and

the following application categories: All Applications, Accessories, Games, Internet, Media, Office and System. Initially, the All Applications choice is selected, and the screen is a mess with the icons for all installed applications on the display.

You can narrow it down by choosing one of the categories, so that only a few icons are displayed. Or, you can click the search field and type part of the application's name or description. This displays not only installed applications that match the search text, but also any matching applications available in the repository. Click the icon of the available program to launch Software Center with that application already selected. If you want the app, just click the Install button.

The result of all this is a display that has saved a significant amount of space and is optimized to use one application at a time. This is how most people work;

you may have other applications running at the same time, but normally you interact with one application at a time.

Unity is available now, in Ubuntu Netbook Edition and as a PPA you can install in the Desktop Edition. (See sidebar for install instructions.) Shuttleworth said, "I'd very much like to get feedback from people trying it out on a Netbook, or even a laptop with a wide screen."

Not all of these features described here are fully implemented, but they should be by the time Ubuntu 11.04 is released. Other things seem unfinished too. I couldn't find a way to make any changes, for example, to customize the launchers on the left panel. Presumably, this will be corrected by the time it's released.

And, if you don't like Unity? Don't worry; you still can install standard GNOME or one of the other desktops from the repository.

Conclusion

Both GNOME Shell and Unity are quite usable today in their current forms. We're not going to see the kind of problems and uproar that occurred when KDE 4 was released. The changes going into the creation of GNOME Shell and even Unity are far less ambitious than the changes that created KDE 4.

Unity also is a shell for GNOME, although it's completely separate from GNOME Shell. Unity and Gnome Shell essentially are opposite ends of the spectrum. Unity is designed for a simple environment where people tend to do one thing at a time, and Gnome Shell is designed for a more complex environment where users are doing multiple activities simultaneously.

Personally, I really like GNOME Shell on a desktop or laptop PC. My impression is that Unity will be a good choice for small touchscreens, but I'd rather have GNOME Shell or even GNOME 2 on a larger display.

But, that's just my preference. If you want something different, your ideal setup is only a few clicks away. ■

Charles Olsen has been working in IT help-desk and technical training for more years than he will admit. He is one of the hosts of *mintCast*, a podcast by the Linux Mint community for all users of Linux. You can find *mintCast* at www.mintcast.org or, if you must, via iTunes.

TRY THEM YOURSELF

GNOME 3 and Unity are not ready for release, but you can get a preview and see where the projects are at the moment. Because they are unfinished, they are not recommended for use in a production environment.

GNOME 3 is easy—just use the package manager to install `gnome-shell`.

Once the package and dependencies have been installed, you can activate it with the command:

```
gnome-shell --replace
```

To install Unity, you need to add a Personal Package Archive. Open a terminal, and enter the command:

```
sudo add-apt-repository ppa:canonical-dx-team/une
```

Enter your password when prompted, and apt will add the PPA and import the key.

Then, update your software list:

```
sudo apt-get update
```

To install Unity, type:

```
sudo apt-get install unity
```

Once Unity is installed, the next time you log in, the interface will default to Ubuntu Netbook Edition. (Ubuntu Desktop Edition still is available, of course.)

The Second-String Desktop

GNOME and KDE may be the heavy-hitters of the desktop world, and although all that power is nice, sometimes it's too bulky. That's where other desktop managers come in.

SHAWN POWERS

It's dangerous for me to use sports metaphors, because my expertise ends at knowing there are three strikes in an out. When it comes to sitting the bench, however, I'm a veteran professional! Most major distributions choose one of the big hitters for their desktop management systems. GNOME and KDE continue the epic battle that keeps the competition intense and our desktops diverse. In honor of this month's Desktop issue, I thought it would be nice to pay some homage to those desktop managers and windows managers that don't get quite as much attention.

Window Managers and Desktop Managers

Before we delve too deep into comparing the various Linux GUIs out there, it's important to understand the difference between window managers and desktop managers. The nuances between the two can be subtle and at times almost nonexistent. A window manager is simply the program running on top of the X server itself that manages windows. Some are very sparse in their features, and some are so robust they approach the usability of a full-blown desktop manager.

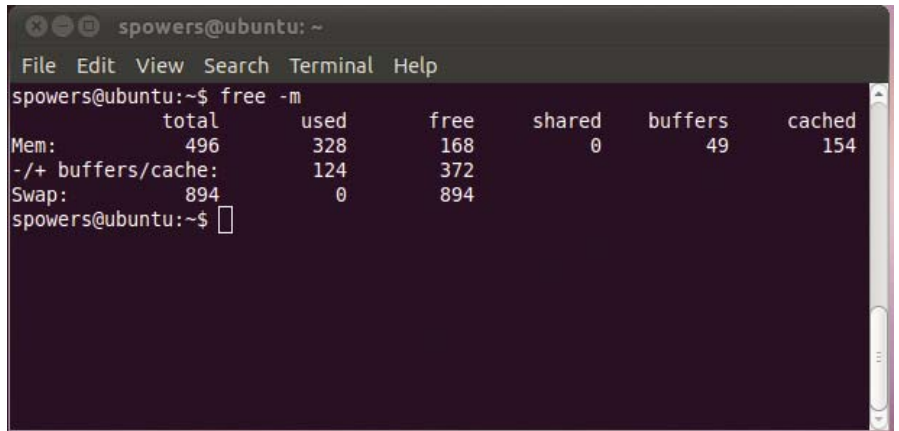
So, what is a desktop manager, you ask? Well, it's more of a total user-interaction interface. It often includes applications, widgets and system integration. In fact, desktop managers (or desktop environments, as they're sometimes called) include a window manager as part of their arsenal. So, although GNOME is a desktop manager, part of the GNOME environment includes Metacity, which is a window manager GNOME uses for, well, managing its windows. It's possible to run a Linux system with only a window manager, as I talk about later, but usually a Linux distribution comes with some sort of desktop manager installed by default.

The Battle for Ubuntu's Bottom End

Ubuntu and Kubuntu certainly are the first-string players for Canonical's Linux lineup. Granted, with Ubuntu 11.04, its flagship product will switch from using a standard GNOME interface to the Unity shell normally used only in its Netbook product, but at least historically, Ubuntu has used GNOME, and Kubuntu has used KDE. Canonical also has its official Xubuntu version for older or less-powerful hardware. Xubuntu runs the XFCE desktop manager. Although it does require fewer resources than GNOME or KDE, many still think XFCE is rather bloated for slower hardware. There is another option, Lubuntu, but it's not officially supported by Canonical. Instead of XFCE, Lubuntu uses the LXDE desktop environment. The Lubuntu team claims to be much less resource-intensive, so I installed them both to see how they "feel" in everyday use.

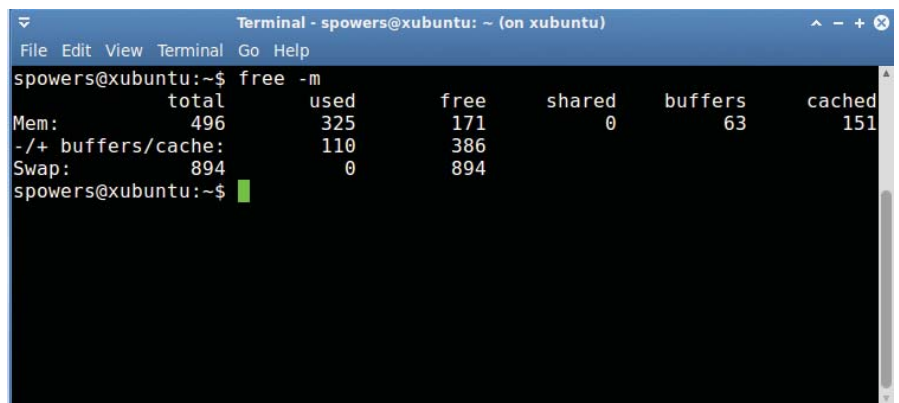
Xubuntu and XFCE

Xubuntu's install is pretty much like every other flavor of Ubuntu. In fact, it's pretty much like every other flavor of Linux. Gone are the days of difficult installs,



```
spowers@ubuntu: ~  
File Edit View Search Terminal Help  
spowers@ubuntu:~$ free -m  
              total        used         free       shared    buffers     cached  
Mem:           496          328          168           0           49          154  
-/+ buffers/cache:      124          372  
Swap:          894           0          894  
spowers@ubuntu:~$
```

Figure 1. A freshly booted, default install of Ubuntu 10.10 uses 328MB of RAM while idle.



```
Terminal - spowers@xubuntu: ~ (on xubuntu)  
File Edit View Terminal Go Help  
spowers@xubuntu:~$ free -m  
              total        used         free       shared    buffers     cached  
Mem:           496          325          171           0           63          151  
-/+ buffers/cache:      110          386  
Swap:          894           0          894  
spowers@xubuntu:~$
```

Figure 2. A freshly booted, default install of Xubuntu 10.10 uses 325MB of RAM—almost identical to Ubuntu.

and even if you choose to use a text-based installer, the process is dead simple—so simple, in fact, it's silly to include a screenshot. It looks like an installer. Trust me.

Current versions of Xubuntu are a little shocking in just how much they resemble their GNOME-y counterparts. In fact, the Xubuntu desktop in version 10.10 looks like a slightly bluer version of Ubuntu 10.10. Certainly this is Canonical's tweaking—a very nice job of making its lighter-on-the-resources desktop look exactly like its big brother. Unfortunately, appearance isn't the only place Xubuntu is identical to Ubuntu. Figure 1 shows a freshly booted install of Ubuntu, with no programs running other than the terminal displayed. You can see the freshly booted new install uses approximately 328MB of the 512MB installed on my machine. When I turned to the Xubuntu install, which runs XFCE instead of GNOME, I expected to see a

much lower memory usage upon booting up. I was shocked to see Xubuntu using 325MB, almost identical to the Ubuntu install (Figure 2).

The big difference with Xubuntu isn't really how much RAM the desktop manager uses, but rather the default applications installed. When I start Xubuntu's Exaile music application versus Ubuntu's Rhythmbox, it does indeed use less RAM, and it starts up faster. However, getting rid of Rhythmbox on Ubuntu and installing Exaile in its place gives the same advantage while using GNOME under Ubuntu. In fact, although Xubuntu and XFCE do feel faster in use, in every case I've tested, it seems to be due only to the default applications. If you're a GNOME fan, Xubuntu might be a big change for little reward. Keep GNOME, install some faster applications, and you might get the best of both worlds.

FEATURE The Second-String Desktop



Figure 3. The Spartan, but Easy-to-Navigate Lubuntu Default Desktop

Lubuntu and LXDE

It turns out I'm not the only person to notice that Xubuntu doesn't really tailor itself to low-end hardware as much as it claims to. The folks in charge of the Lubuntu Project decided improving performance for slower machines should include more than installing zippy applications by default. Just like with Xubuntu, Lubuntu installs in an extremely unexciting and completely functional way. Once installed, however, Lubuntu does differ from Xubuntu and Ubuntu in appearance.

Although a similar blue to Xubuntu, Lubuntu's screen layout is visually different. Figure 3 shows Lubuntu's simple single taskbar layout, which is quite similar to the design Microsoft has been trying to perfect since the days of Windows 95. That's not a bad thing. As a group we may not care for Microsoft, but its start-menu-type system is widely known and very usable. The first thing I did upon booting Lubuntu was open a terminal and check the memory usage. You can see in Figure 4 that Lubuntu is using only 163MB of RAM when fully booted. That's almost exactly half the RAM Xubuntu and Ubuntu use when freshly booted. When you add to that the selection of fast and small applications Lubuntu uses by default, including the Chromium Web browser, it really does scream even on low-end or old computers. If you've been frustrated with your computer's poor performance with either GNOME under Ubuntu or XFCE under Xubuntu, you may want to give Lubuntu a try. It uses the same repositories, and it has remarkable compatibility with more well-known, if not bulkier, applications.

What about Others?

There are more to the alternatives than just speed. I used the Ubuntu example

```
spowers@lubuntu:~$ free -m
              total        used         free       shared    buffers     cached
Mem:           496         163          333           0          13          88
-/+ buffers/cache:
Swap:          894           0          894
```

Figure 4. A freshly booted, default install of Lubuntu 10.10 uses only 163MB of RAM while idle.

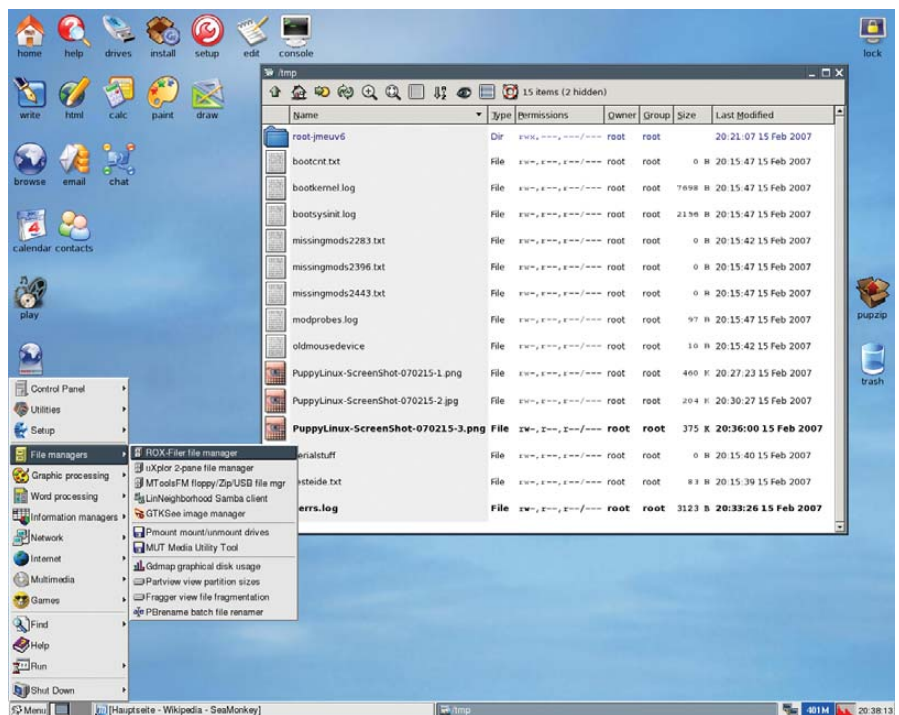


Figure 5. Puppy Linux Using ROX-Filer (screenshot by Joachim Köhle)

above to demonstrate how alternatives to the "Big 2" can give you advantages in speed, but there are other reasons to veer from the norm as well. The ROX Desktop, for example, is a complete desktop environment designed around a file manager—the ROX-Filer, to be specific. Although the ROX Desktop is certainly light on system resource needs, its design and integration with

the filesystem is what really makes it unique. Figure 5 shows a screenshot of Puppy Linux, which uses ROX-Filer as the file manager.

The ROX Desktop suite includes its own window manager, OroboROX, but like most other desktop environments, it doesn't rely on one specific window manager to work. When you find a window manager you like, it's often

Table 1. A Sampling of Linux Desktops/Window Managers

DESKTOP/WINDOW MANAGER	DESCRIPTION	DESIGN GOALS	BASED ON	ADVANTAGES	DISADVANTAGES
KDE	Full desktop environment	Full system integration, including applications	Uses KWin window manager and Qt libraries	Great application integration, highly customizable	Distinct look; non-KDE apps often seem awkward
GNOME	Full desktop environment	Full system integration, including applications	Uses Metacity window manager, based on GTK+ libraries	Wide variety of native applications, wide adoption in corporate environments	Non-GTK apps often look odd and use more RAM
LXDE	Lightweight desktop environment	Speed and beautiful interface	Uses Openbox window manager and GTK+ libraries	Works well on older/slower hardware, maintains compatibility	Lacks some of the features found in GNOME or KDE
XFCE	Lightweight desktop environment	Full-featured desktop environment, but light on resources	Usually uses XFWM4, but works well with other window managers	Somewhat lower system requirements than GNOME or KDE	Possibly a bit too resource-hungry for low-end systems
Enlightenment E17	Window manager with the features of a desktop manager	Speed and eye candy with integrated functionality	A window manager plus a set of libraries for developing apps	Fast without sacrificing style	Still in beta but quite stable
ROX Desktop	Desktop manager based on the ROX-Filer	Approaches the OS in a file-centric way	ROX-Filer file manager and the OroboBox window manager	Unique file-based design makes installing apps drag and drop	ROX Desktop is either a love or hate affair
IceWM	Hybrid window manager and desktop manager	Speed and simplicity	Simple menu and taskbar design	Fast and easy to make system-wide configuration changes	No way to make desktop icons, requires additional software for some features
Blackbox/Fluxbox	Very minimalistic window managers	Speed and small memory/CPU footprint	Fluxbox is based on Blackbox (it's a fork)	Blazingly fast	Very limited in features, but by design not immaturity
Openbox	Very minimalistic window manager	Speed and small memory/CPU footprint	Originally based on Blackbox, original code since version 3.0	Simple and fast	Limited in features by design
AfterStep/Window Maker	Clones of the NeXTSTEP interface	Functions and looks like NeXTSTEP	Designed after the unique design of the NeXTSTEP interface	Unique	Often difficult to configure, and the interface is an acquired taste
Ratpoison	A window manager that doesn't require a mouse	Kills the need for a mouse	Designed after GNU Screen	No need for a mouse	Very limited in features, which the developers consider a feature
DWM	An extremely minimalist window manager	Manages windows and nothing more	The ideas of other minimalist window managers	Small and fast	No configuration files, must edit source code to configure

possible to use it seamlessly with whatever desktop management system you want. In fact, many people, and even entire distributions, run only a window manager. This is possible because many window managers are so feature-rich, they do most of the things a desktop manager would do. One good example of that is Enlightenment.

The Enlightened Don't Need a Desktop Manager

Enlightenment is a window manager that has been around a long time. Some window managers are minimalistic; however, Enlightenment is extremely feature-rich. It provides a file manager, a dock, a GUI configuration tool, application launchers—pretty much everything required in a

full-blown desktop environment. Does that mean it's a desktop manager and not just a window manager? Perhaps. It doesn't really matter how you define it though. Enlightenment is one of those things everyone should try at least once. There are even live CDs, specifically designed for trying Enlightenment. Figure 6 shows applications running under the Elive CD

FEATURE The Second-String Desktop

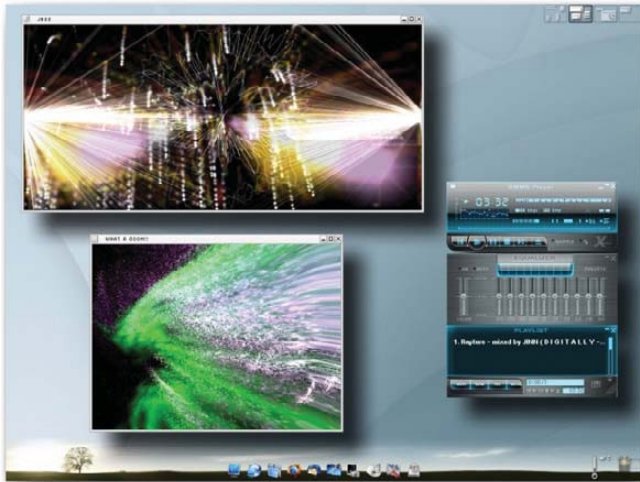


Figure 6. Enlightenment E17 Desktop (from the elivecd.org Live CD)

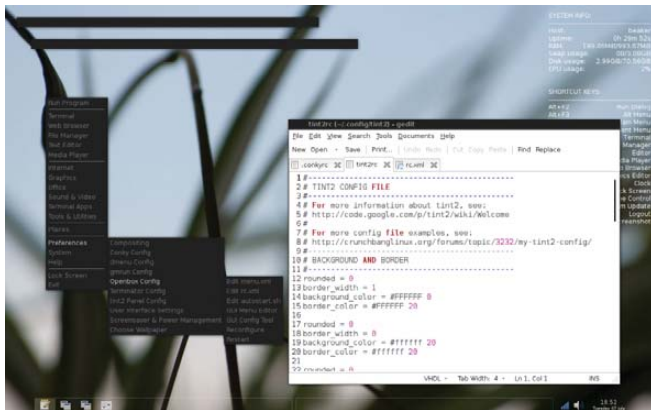


Figure 7. CrunchBang Linux is a minimalistic distribution that uses Openbox as the window manager.

default desktop.

IceWM is another window manager that is rather profound in the features it offers. Although it doesn't have a mechanism for creating desktop icons, it does have a very robust menu system and application suite for managing most aspects of the Linux desktop. IceWM is very customizable, and although it uses the familiar Windows-like start menu, it doesn't try to clone Microsoft. In fact, I use a combination of IceWM and Nautilus on my network of 150 older thin clients because it's fast and reliable. Because the menu system is controlled with a single system-wide config file, it makes wide-scale customization a breeze.

A multitude of Linux distributions have only a window manager to manipulate the desktop. Blackbox, Fluxbox and Openbox are all related window managers. Fluxbox is a fork of Blackbox, and although Openbox is all original code

now, it started as a fork of Blackbox as well. These three window managers are lightning fast. They may not offer the same level of features and complexity that Enlightenment or IceWM do, but for many minimalistic distributions, they are just perfect. CrunchBang Linux is a prime example of a full-featured, yet minimalistic distribution. It uses Openbox as its window manager, and as you can see in Figure 7, the windowing environment is designed to get out of the user's way.

Choice, the Ultimate Freedom

The great thing about choosing a Linux GUI is that it costs you nothing to change. Whether you like Fedora's default GNOME install

or prefer openSUSE's green-lizard KDE install, every Linux install can be tweaked or changed. I must warn you though, once you try the NeXTSTEP clones Window Maker or AfterStep, you might never want to see a start menu again. If you experiment with the mouseless beauty of the Ratpoison window manager, you might never want to click again. Or, perhaps you're at the other end of the spectrum, and you want to fool yourself or others into thinking you are running OS X. Check out the free Macbuntu GNOME theme. With Linux, customization is king, and the options are seemingly endless. I created a chart to help you sort out some of the Linux GUI options available (see Table 1). It's by no means an exhaustive list, but it should get you started. Remember, just because a desktop environment is sitting on the bench doesn't mean it didn't make the team. Check out the bench-warmers, you might just find a winner.■

Shawn Powers is the Associate Editor for *Linux Journal*. He's also the Gadget Guy for LinuxJournal.com, and he has an interesting collection of vintage Garfield coffee mugs. Don't let his silly hairdo fool you, he's a pretty ordinary guy and can be reached via e-mail at shawn@linuxjournal.com. Or, swing by the #linuxjournal IRC channel on Freenode.net.

Resources

AfterStep: www.afterstep.org

Blackbox:
blackboxwm.sourceforge.net

CrunchBang Linux:
www.crunchbanglinux.org

DWM: dwm.suckless.org

Elive: www.elivecd.org

Enlightenment E17:
www.enlightenment.org

Fluxbox: www.fluxbox.org

GNOME: www.gnome.org

IceWM: www.icewm.org

KDE: www.kde.org

Lubuntu: www.lubuntu.net

LXDE: www.lxde.org

Macbuntu:
macbuntu.sourceforge.net

Openbox: www.openbox.org

Puppy Linux: www.puppylinux.org

Ratpoison:
www.nongnu.org/ratpoison

ROX Desktop:
rosidus.com/desktop

Ubuntu: www.ubuntu.com

Window Maker:
www.windowmaker.org

XFCE: www.xfce.org

Xubuntu: www.xubuntu.org



DRUPALCON CHICAGO

MARCH 7-10 2011

Drupal is the industry-leading open source content management platform that's used to power millions of websites around the world.


This March 7-10, thousands of Web developers, designers, businesspeople, and everyday citizens will gather for DrupalCon Chicago, a one-of-a-kind conference that no Web professional can afford to miss!

Find out more at <http://chicago2011.drupal.org/>



MARCH 7-10, 2011 ⚙ **SHERATON HOTEL & TOWERS** ⚙ **CHICAGO, ILLINOIS**

Organize Your Life with Nepomuk



Can KDE bring order from chaos? **Stuart Jarvis**

When you build a house, you cannot start with the paint, fancy windows and doorbell. Instead, you spend a lot of time digging in the ground, disturbing the neighbors and making little visible progress while you lay the foundation. It has been much the same with Nepomuk, the Semantic Desktop technology of the KDE platform. Work has been going on for a long time beneath the surface, causing occasional disruption—such as file indexing slowing down the rest of your desktop—with little visible progress. However, the foundations now are solid, the main structures are in place, and KDE's developers are adding features that make Nepomuk useful for you, right now.

So, what can Nepomuk actually do for you?

Stop Searching, Start Finding

The first thing Nepomuk offers is file searching integrated into your KDE applications. If your distribution has enabled Nepomuk, files also can be tagged and rated from KDE's Dolphin file manager. You can search for files simply by typing a query into Dolphin's search box, which brings up results and a panel of basic options that make it easy to refine the search. These enable you to limit searching to the current folder and let you choose whether to search everything, text documents (including OpenDocument texts and PDFs), images or filenames. When you have identified the files you want, you can use the Save button to add the search query to your Places panel on the left of the Dolphin window, making the search available in the future via a single click. The saved search also is available in the file chooser dialog of all KDE applications.

You can access additional search options in Dolphin 1.5 (part of KDE software compilation 4.5) by clicking the green + button. These provide filtering based on file modification date, size, tags or rating. You can continue to add and combine additional filters until you have isolated the exact files you want—for example, by limiting the search to files with a particular tag that have not been modified since a particular date. You also can use the filters directly without any search terms.

When Nepomuk first came to KDE software, it lacked a good graphical search interface to expose its capabilities. Nepomuk always has offered the ability to

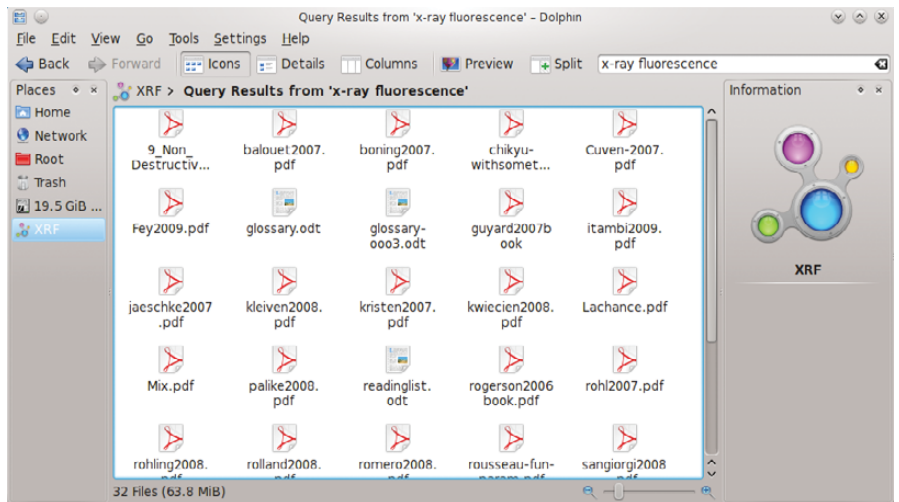


Figure 1. You can access saved searches from the Places sidebar.

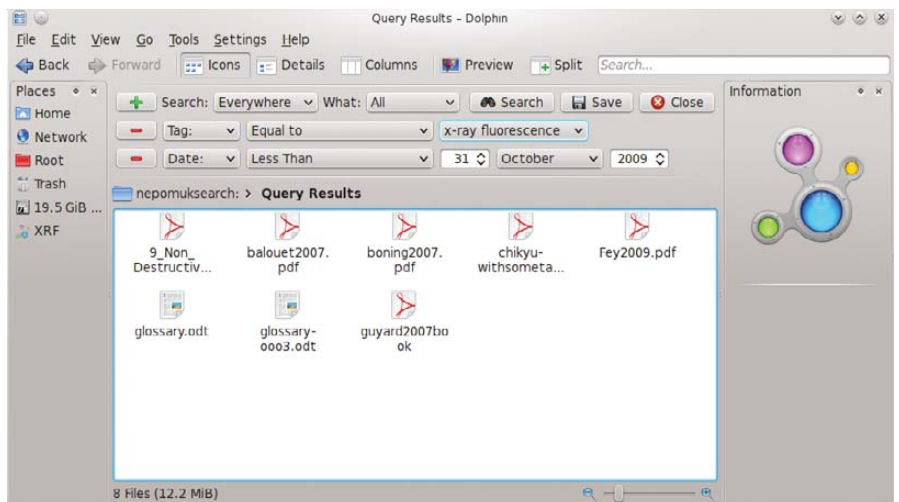


Figure 2. Dolphin 1.5 lets you combine filters to narrow your search.

What Is This Thing Called Nepomuk?

Nepomuk (also known as the Networked Environment for Personalized, Ontology-based Management of Unified Knowledge) began as a research project funded by the European Union to explore better ways of managing, understanding and sharing desktop information. It attracted participation from corporations, such as IBM and Linux vendor Mandriva, in addition to involvement from the National University of Ireland.

The initial project ran from 2006 to 2008, during which time a reference implementation was written in Java. Mandriva and KDE have since worked on integrating the ideas into KDE software, retaining the name Nepomuk.

Practical use of the technology required a high-performance file indexer and query database. Initially, Nepomuk's own Java-based Sesame query framework was used in the KDE implementation, but its reliance on Java led to packaging problems for some distributions, and KDE switched to Virtuoso in 2009. No longer held back by these technical constraints, the main barriers to making greater use of Nepomuk in KDE software have been removed.

construct complex search queries using a query language, such as SPARQL. You could, and still can, enter queries in this way using Dolphin. Click on the breadcrumbs above the search results to access Dolphin's editable location bar, view the query that was used to generate the results and edit it directly. However, many users do not have an intimate knowledge of the query language constructs and would consider such an approach better suited to the command line than a graphical file manager. The search interface in Dolphin 1.5 graphically exposed many of Nepomuk's search capabilities for the first time.

Peter Penz, Dolphin's lead developer, still sees problems with the search interface in Dolphin 1.5: "It just takes way too many clicks to specify a query." KDE

FEATURE Organize Your Life with Nepomuk

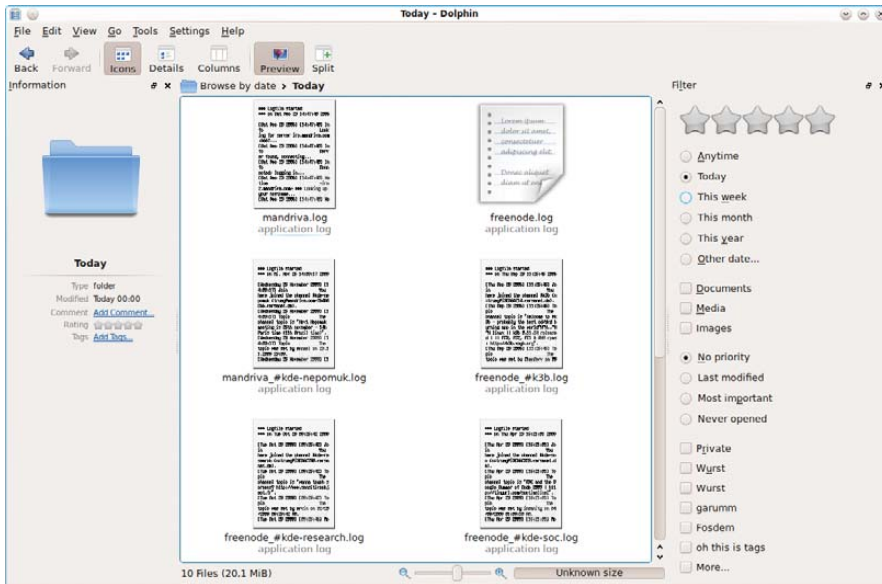


Figure 3. Faceted browsing makes it easier to combine filters in Dolphin 1.6.

Although computers are very good at indexing files, they are very bad at understanding the relationships between them.

developers, including Peter and Mandriva's Sebastian Trüg have been grappling with the problem of making a good graphical interface that exposes the power of Nepomuk searches in a more convenient manner. In 2009, Alessandro Sivieri began work on "faceted browsing", an approach to the problem that provides panels of search filters in the file manager sidebar, resulting in his Sembrowser prototype (see Resources). Now, these ideas have come together to provide a faceted browsing sidebar in Dolphin 1.6 (part of KDE software compilation 4.6) that appears when you start a search and makes additional filters available with single clicks. Faceted browsing also will be integrated into the KDE platform, making it available for other KDE applications to use.

To Index or Not to Index?

Nepomuk's text searches rely on your computer having an index of files and their contents. This is created by the Strigi indexer, which claims to be the "fastest and smallest desktop searching program". That may be true, but crawling and interpreting files always is going to be a

resource-intensive task. KDE mitigates this by indexing files only when the processor




Figure 4. File indexing can be suspended with a couple clicks.

load is otherwise low and suspending indexing when you switch to battery power. You also can suspend (or resume) indexing on demand using Nepomuk's system tray icon, which is visible whenever the indexing service is active. To suspend the indexer manually, just right-click on the tray icon and check Suspend File Indexing. In my experience, it is much more likely that you will find the indexer has been suspended automatically at times when you would not mind having it running than you will need to suspend it manually.

To gain full control over indexing, simply visit the Desktop Search page in KDE's System Settings configuration utility. Alternatively, you can access the configuration options directly from the right-click

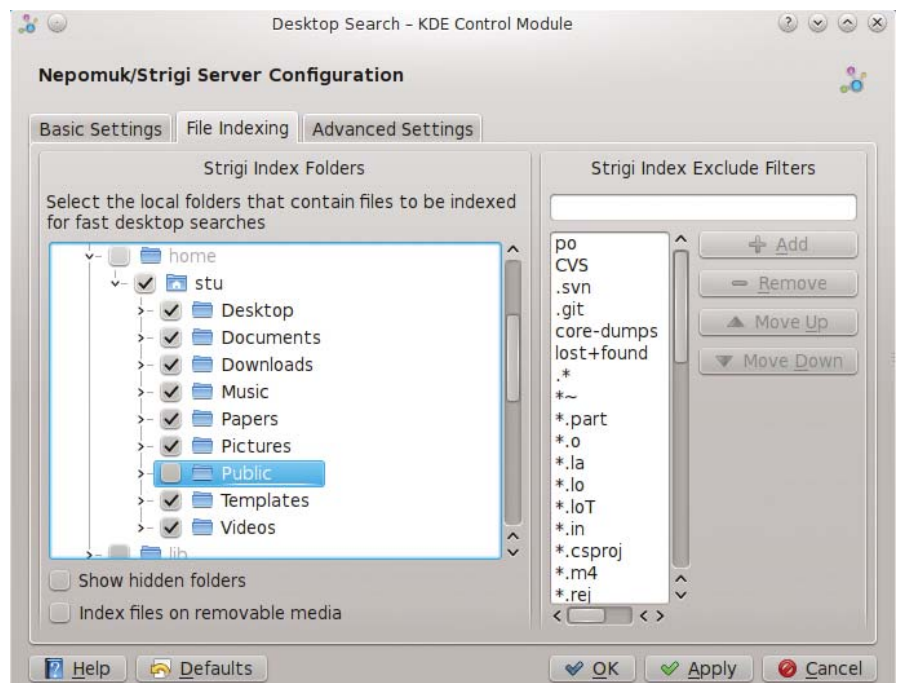


Figure 5. KDE System Settings gives you fine control over file indexing.

menu of the Nepomuk system tray icon. You have the choice of whether to enable Nepomuk and the Strigi file indexer at all, and you can specify which folders should be indexed. You also can exclude certain file types, such as backup files and partial downloads, and even set a limit on the memory used by the search database.

Once the file indexer has collected data for all of your files, it has to be able to search the data quickly and efficiently. This was a major problem facing Nepomuk when it first was included in KDE software, as the available databases either lacked performance, features or both. However, since the release of KDE software compilation 4.4 in February 2010, Nepomuk has used the Virtuoso database, and as Peter puts it, "Performance problems are nearly gone."

Still, it makes sense not to index everything. After all, you probably won't often need to find one of the thousands of shared library or application files on your computer, and including them would waste resources.

However, you occasionally might need to find one of those files. Before Dolphin 1.6, terms entered in Dolphin's search bar queried only the Nepomuk database, and it was necessary to use the separate KFind interface to search non-indexed files. This meant you had to know whether the file you were looking for was in an indexed location before deciding which search interface to use. In Dolphin 1.6, both search types are integrated in a single search box, so that the Nepomuk database is used for directories that have been indexed by Strigi, and KFind is used for other locations. In this way, you still get search results when you search in a folder that is excluded from indexing, it just takes a little longer.

You Are Still Smarter than Your PC

Fast and customizable desktop search is the most visible and, at present, probably the most-used benefit provided by Nepomuk in a KDE workspace. However,

the Nepomuk developers have much higher ambitions. Although computers are very good at indexing files, they are very bad at understanding the relationships between them.

Making meaningful relationships between files remains one of the major challenges for Nepomuk. If your friend Alice is getting married and you need to find the map you were sent that shows the wedding venue's location, you might look for e-mail messages from both her and her future husband, Bob. You instinctively would realize that information related to Alice may be found in a Web page or communication from Bob, because you understand their relationship. However, your computer would not give any special priority to files that had come from Bob.

One way of making your computer act a little smarter is to tag your files. This works in a similar way to how you tag music files, so your media player can present you with albums, artists and even genres to browse, instead of forcing you to navigate through directory structures.



visit us at www.siliconmechanics.com
or call us toll free at 866-352-1173



**Powerful.
Intelligent.**



Charles heads the web development team here at Silicon Mechanics: he's responsible for the configurators and power calculator on our site, among other things. As a software expert, he offers a useful perspective on our server and storage products.

When asked what he would do if he had a Rackform iServ R413 configured with 4 8-core Intel® Xeon® Processor 7500 Series CPUs and 32 DDR3 DIMMs, he said, "32 virtual machines . . . one per core . . . in one rack unit." But he didn't stop there.

Charles knows that to make the best use of a server with that kind of processing horsepower in a virtualized environment, he needs I/O to match. He paired the 4P server with a Storform iServ R516 storage server, configured with 24 2.5-inch Intel X25-E solid state drives. Think of it as a developer's dream team: multi-core processing and high memory counts for blistering performance, and high-performance storage for blazing I/O speed.

Need a "dream team" of your own? Talk to the Experts at Silicon Mechanics for the perfect match.

When you partner with Silicon Mechanics, you get more than just the power and performance of the latest Intel technologies—you get an Expert like Charles.



For more information about the
Rackform iServ R413
visit www.siliconmechanics.com/R413

Expert included.

Silicon Mechanics and the Silicon Mechanics logo are registered trademarks of Silicon Mechanics, Inc. Intel, the Intel logo, Xeon, and Xeon Inside, are trademarks or registered trademarks of Intel Corporation in the US and other countries.

FEATURE Organize Your Life with Nepomuk




Figure 6. The Gwenview image viewer makes it easy to tag your photos.

If going through all of your files and manually adding tags and ratings is not your idea of fun, good news is coming for you.

You can tag any kind of file simply by clicking on the Tags section of Dolphin's Information panel. This brings up a dialog to apply existing tags or create new ones. KDE image viewer Gwenview also enables you to use and edit the same tags and to browse images by tag rather than path. If you received a map from Bob relating to his wedding to your friend Alice, you could tag it with terms such as "Bob", "Alice", "wedding" and "map" and then easily find it using Dolphin's search box.

There are other possibilities for tagging beyond making search more powerful. If you are working on a number of projects and some of your files are relevant to more than one of those projects, you can apply tags for each project, filter based on those tags in Dolphin, and then save that search as a "Place". You then can access the files relevant to a given project in any KDE application quickly, without needing to copy any of the files into project-specific folders or set up links.

Of course, tagging your files manually could take a lot of time, and it is

not information you would want to lose when you buy a new computer or re-install your operating system. With KDE platform

4.6, Nepomuk provides the ability to back up and transfer this human-generated metadata so you can always take it with you.

You also can use tagging and saved Nepomuk searches to have a selection of relevant files available right on your desktop when you log in. To do this, all you need to do is create a Folder View Plasma widget in your KDE workspace, and then use its configuration dialog either to enter a search query directly or select a search you already saved as a Place. For example, you could have folder views containing the files for each of the main projects you are working on. You even could restrict the initial view to only the most important files by rating them with five stars and filtering the results by rating.

It's Good to Share

Tagging is, of course, in itself, nothing new. Media players and photo management applications have been using tags for years to enable convenient views of files that are freed from a single, rigid directory structure. However, these systems have been largely incompatible. You can read basic tags embedded in Vorbis music files in KDE's Amarok and in other music players, but you cannot easily transfer other metadata, such as ratings or playcounts, between music players. Neither do tags applied to images in KDE's photo management application, digiKam, appear in Dolphin

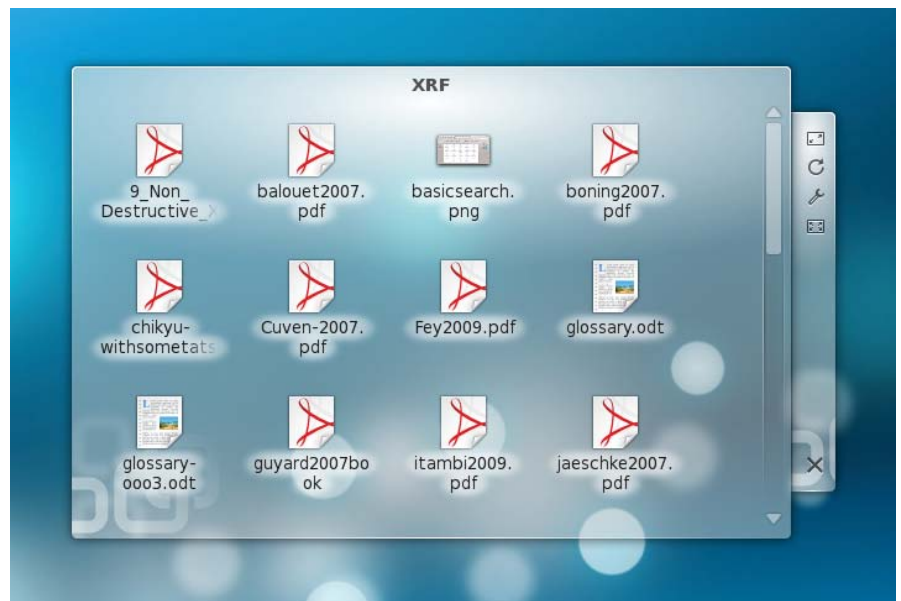


Figure 7. You can use saved Nepomuk searches to put relevant files on your desktop.

or Gwenview. For some file types, it is not possible to embed metadata at all.

As more applications make use of tagging to simplify organization and retrieval of content, there is a danger of redundancy with the same data being collected and stored multiple times, so sharing storage and resources makes sense. To work together, applications not only need to have a way of sharing metadata but also a way of understanding it. One of the main jobs of the Nepomuk Project is to define these common ways of representing metadata of differing types. E-mail messages that came from your friend Alice are different from e-mail messages that mention her, and you want your computer to understand that difference.

Within the world of KDE software, the most visible users of Nepomuk are the Dolphin file manager, Gwenview image viewer and the Bangarang media player. The last integrates browsing of media by type (movies, TV shows, short clips), rating, playcount and date of last play. Although use of Nepomuk in KDE applications is limited at present, Peter is optimistic about the future take up of Nepomuk due to improvements in the ease of accessing the technology: “searching and filtering—it gets easier and easier from applications in each release”.

Be Lazy

If going through all of your files and manually adding tags and ratings is not your idea of fun, good news is coming for you.

Nepomuk developers are working on automatically identifying the files you use most often and prioritizing those in search results. File use also can be linked to Plasma Activities, so that you could be presented with different files depending whether you were using your computer for work or play. Nepomuk also is able to keep track of the source of files you find on-line, recording the original file location. This not only helps you find images you have downloaded from your friend’s on-line photo album, but also could work the other way around, enabling you to find the Web site where you downloaded a favorite desktop background.

KDE’s personal information and communication suite, Kontact, is being ported to a new storage system, Akonadi (see Resources). This will make it easier to maintain semantic information on

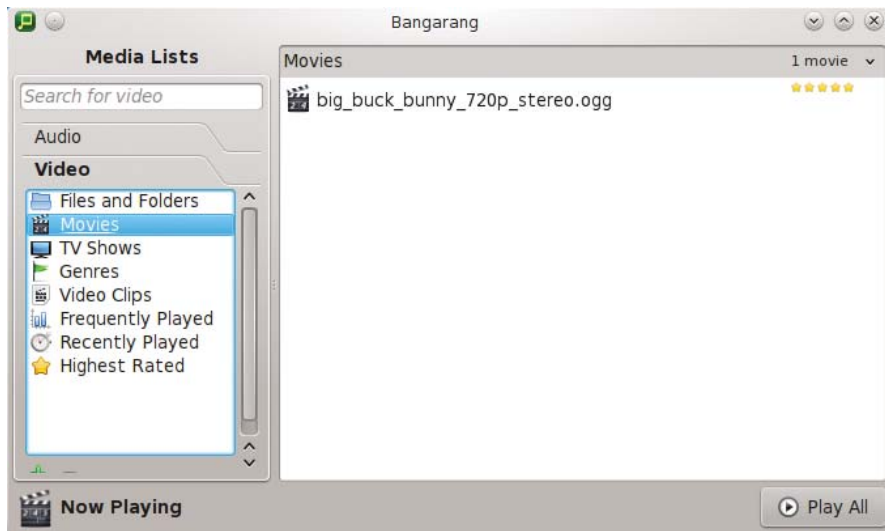


Figure 8. Bangarang lets you browse using information from Nepomuk.

items that arrive in messages—for example, by linking file attachments with the person who sent them—so that you can find an image your friend sent you two weeks ago.

KDE and Mandriva developers also are working on integrating technologies from Scribo (Semi-automatic and Collaborative Retrieval of Information Based on Ontologies). This aims to automate the extraction of meaningful metadata from files. For example, it should be possible to determine keywords and topics from text-based documents to suggest tags or even apply them automatically. Scribo also has the ability to extract text from images, so that a photograph including a road sign bearing the word “Paris” might be found in a search for “Paris” without the need for manual tagging. Capabilities like this are not yet integrated into the KDE platform, but experimental code has been included in the KDE source repositories.

If you are curious about possible future directions of Nepomuk, Mandriva Linux 2010 includes some experimental features. It has special versions of KDE’s Konqueror Web browser and KMail messaging program that enable tagging of Web pages and e-mail so they can be retrieved in searches. It also has the ability to offer tag autocompletion—for example, suggesting the tag “John Smith” from your address book when you begin to type the word “John” as a new tag.

This Is Just the Beginning

It would be wrong to describe Nepomuk as feature-complete or without any flaws. Its developers see much greater potential than that currently realized in KDE software. However, enabling Nepomuk and Strigi indexing in your KDE workspace today gives you unobtrusive, fast and flexible desktop search to find your files easily. You also can begin to organize your files using tags rather than being limited to rigid directory structures, and you can save search queries as Places for easy access from any KDE application. You even can have sets of your most important files waiting for you when you log in. The Nepomuk house is not complete, but the foundations are strong, and it is now safe to go in and take a good look around. If you do so, you may discover you like what you find. ■

Stuart Jarvis is a scientist and member of KDE’s Marketing Working Group. He divides his time between losing data files, graphs and papers and finding them again.

Resources

Sembrowser: kde-apps.org/content/show.php?content=117692

Strigi: strigi.sourceforge.net

Akonadi: pim.kde.org/akonadi

Scribo: www.scribo.ws

Create Your Own Linux Desktop *and Take It with You*

Configure your own custom desktop and plug it in to any PC to re-create your working environment wherever you go.

RICK ROGERS

Most of us work with multiple Linux machines these days, and we are used to hopping between them. However, there are times when you want to separate your Linux desktop experience from your physical desktop PC. You would like a Linux desktop that lives on a USB key that you can move freely among PCs. To explain this better, let me describe two problems I've recently addressed with a small, portable Linux distro:

1. I don't always work from the same location. Carrying a portable laptop computer around is one way of having my Linux desktop wherever I need it, but another way is to have a small Linux desktop on my keyring. When I arrive at a new work location, I can plug the USB key in to any available PC, boot Linux quickly and continue working right where I left off. The live Linux on the key runs entirely in memory and uses the key for persistent storage. I don't need a lot of storage (most of my data lives in the cloud), but I do need to be able to use a unique set of application packages so I have the tools I need to do my work.
2. I have a Netbook that bit the dust a few months ago. The failure was in the hard disk, and this particular generation of Netbook uses a 1.6" hard disk. If you can find a replacement drive that size, the price is such that it's cheaper to purchase a new Netbook. I'd like to be able to run my Linux desktop even on this machine, which is otherwise a brick.

To address these problems, I needed a Linux distro with the following characteristics:

- Small enough to fit on a USB key.
- Bootable on a variety of machines (all PC architectures).
- Ability to run in memory only, given target machines would all have 1GB or more.
- Quick booting.
- Application package compatibility with a major Linux distro.

So, why not just run a desktop distro from a USB stick? That, in fact, was my first solution, and most desktop distros make it really easy to create a live USB installation, including persistence to storage on the USB drive. For example, Ubuntu 10.10 (and every Ubuntu release since 9.04) comes with the `usb-creator` tool that creates a live USB installation for you given the related `.iso` file.

To create an Ubuntu 10.10 live USB installation, first download the `.iso` file for the desktop version from the Ubuntu Web site (www.ubuntu.com/desktop/get-ubuntu/download). After the download, plug in a USB stick (2GB or larger), and select System→Administration→Startup Disk Creator from the Ubuntu

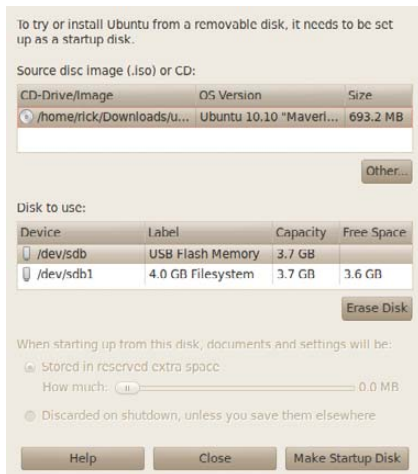


Figure 1. Make Startup Disk

menu. Select the downloaded .iso and the target USB device—in the case shown here, it's /dev/sdb1. When you click Make Startup Disk, Ubuntu copies all the needed files to the stick and makes it bootable.

You'll get a confirmation box, and then you've got your portable Ubuntu. If that meets your needs, you're all set. You can boot it on another machine, install applications using `sudo apt-get install`, and do pretty much anything you would do on your desktop, limited only by the available storage on the USB stick.

Smaller, More Custom Distros

But, you may have needs that go beyond what you get with a standard off-the-shelf desktop. Because this is *Linux Journal*, you almost certainly have custom needs and preferences, and fortunately, there are ways to create your very own portable desktop distribution.

Puppy Linux, in particular, is built with customization in mind. If you've never heard of it, you may want to take a look at Louis Iacona's article "Puppy Linux" in the April 2008 issue (www.linuxjournal.com/article/9932). You also should visit www.puppylinux.com. Barry Kauler created Puppy Linux in 2003, and he has continued to develop it and build a community of developers for it since then. Lately, Barry has focused more on customization capabilities for the platform, resulting in the Woof tool, which builds Puppy-like distributions.

Puppy Linux has many other desirable features, including:

- Runs entirely from memory, so it's very fast.
- Very small memory footprint (100–110MB).
- Support for a wide variety of hardware devices.
- Compatibility with packages from other distributions (including Ubuntu).
- User interface that's intended to be approachable.

There also are some features that might not be desirable, depending on your application:

- Runs entirely in memory, so all active programs have to fit.
- User runs as root, so there's the potential to screw up.
- Hardware device support not as complete as major distros.

Even with the tools that Puppy Linux provides, building your own distribution isn't for the novice. You need to be comfortable using the Linux command line, have a suitable build environment and a broadband Internet connection. Additionally the tools still are under development, so expect that things are not entirely complete (such as reports and so on), and that some combinations of configuration choices just may not work yet. The best news is there's a very active

community of Puppy and Woof users you can reach for help at puppylinux.org.

Building your distribution proceeds in five stages, with each stage having a number of steps:

1. Create a Puppy Linux host system. The Woof tool works only when run under Puppy Linux, so the first thing you need to do is create a host system that runs PL.
2. Use your PL host to download the Woof tool.
3. Use the `woof_gui` tool to configure the files that Woof will use to build your distribution.
4. Build your distribution.
5. Install your distribution on a bootable USB key.

Stage 1: Create a Puppy Linux Host System

Barry and the PL community have made this stage extremely easy. Simply download the latest Puppy Linux .iso file (LuPu 5.1.1 at the time of this writing) from one of the repositories listed on puppylinux.org, and write the .iso to a bootable CD. On Ubuntu, you can use System→Administration→Startup Disk Creator to burn the CD. For Windows, the PL site recommends the BurnCDCC utility.

Once you have your PL CD, you need to boot it on a Linux system. Puppy Linux runs completely from memory, but Woof needs 10GB of file space to do the build,



Figure 2. bones Setup

Puppy Linux, in particular, is built with customization in mind.

so you'll want to use the underlying hard disk. Unfortunately, Woof works only with an ext3 or ext4 filesystem, so a Windows disk won't work for you. When PL boots, it should display a desktop and allow you to start the SNS utility to configure your network connection.

At this point, you could use the Puppy Linux Install utility (on the desktop) to create a PL USB key. If you like PL as it comes preconfigured, or just want to install a few additional applications from .deb packages, you can stop here, without going through the full build of a distribution. You also could install Puppy Linux to your hard disk if you want, but you don't have to.

Stage 2: Download the Woof Tool

PL comes with its own version control system called Bones. The first thing you want to do is create a directory called woof-tree somewhere on your hard disk. In the setup I used, PL sees my Linux hard disk as /dev/sda1.

Now you can open a PL terminal window (Menu→Utility→Urxvt terminal emulator), and note that you're running as root (#). In my setup, I cd'd to my Linux home directory and created woof-tree:

```
# cd /mnt/sda1/home/rick
# mkdir woof-tree
# cd woof-tree
```

Once that's done, you first tell bones which repository:

```
# bones setup
```

You can fill in anything you like for local_username; the Download URL is bkhome.org/bones/woof.

Now, enter:

```
# bones download
```

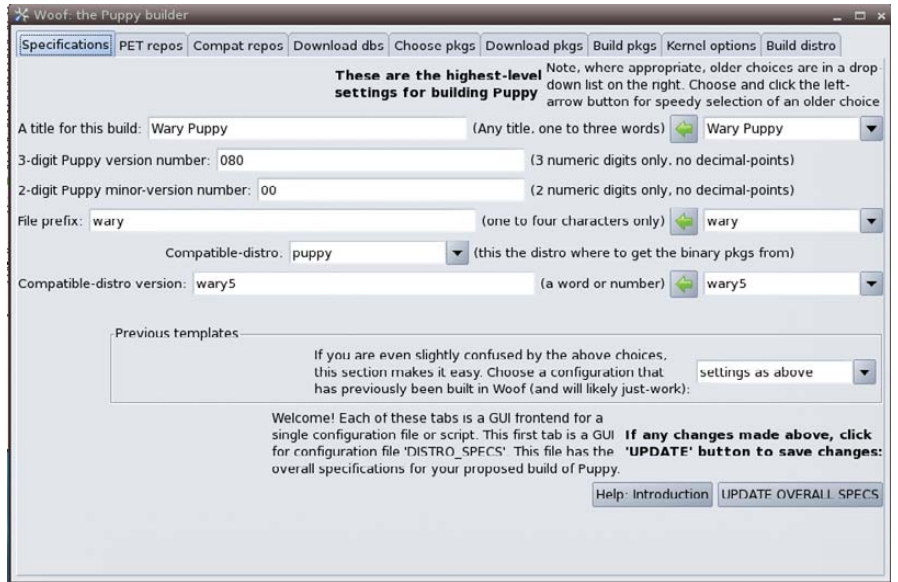


Figure 3. Woof GUI Spec Wary

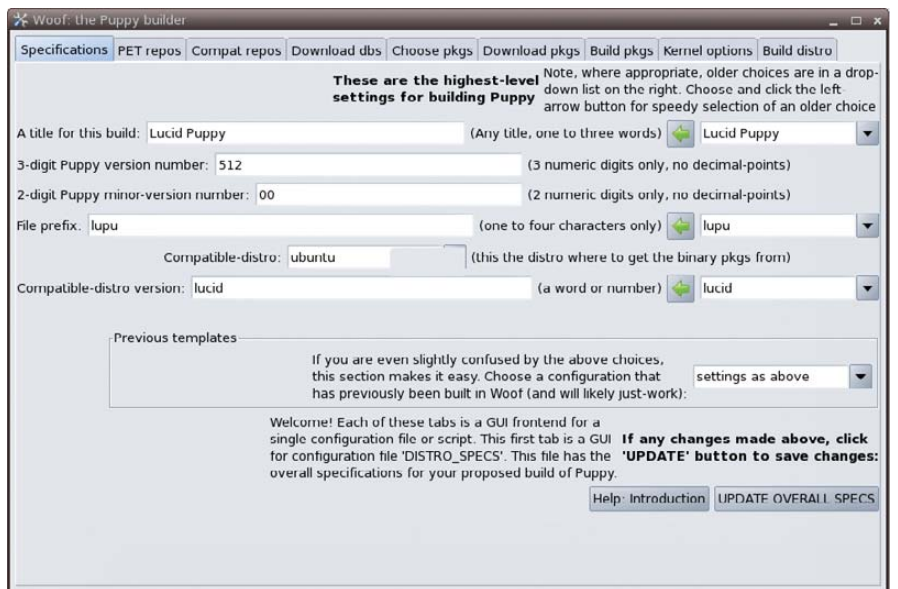


Figure 4. Woof GUI Spec Lucid

And, bones will do just that—download and unpack the current set of Woof files.

Once that's done, it's a good time to read the downloaded file README.txt, which you'll find in woof-tree. This file contains the latest information about the tools and also serves as the help file for the next stage.

Stage 3: Configure Woof

Bones downloaded a number of files into woof-tree, including some configuration files with names beginning in DISTRO_,

some other files and scripts, and six numbered scripts: 0pre, 0setup, 1download, 2createpackages, 3bulddistro and 4quirkybuild, which isn't used for Puppy Linux builds.

You can edit the configuration files and use Woof from the command line to build your distribution, but I'll warn you that it's a bit daunting. Luckily, Barry and friends have provided a GUI interface to the build system that helps you through the process:

```
# ./woof_gui
```

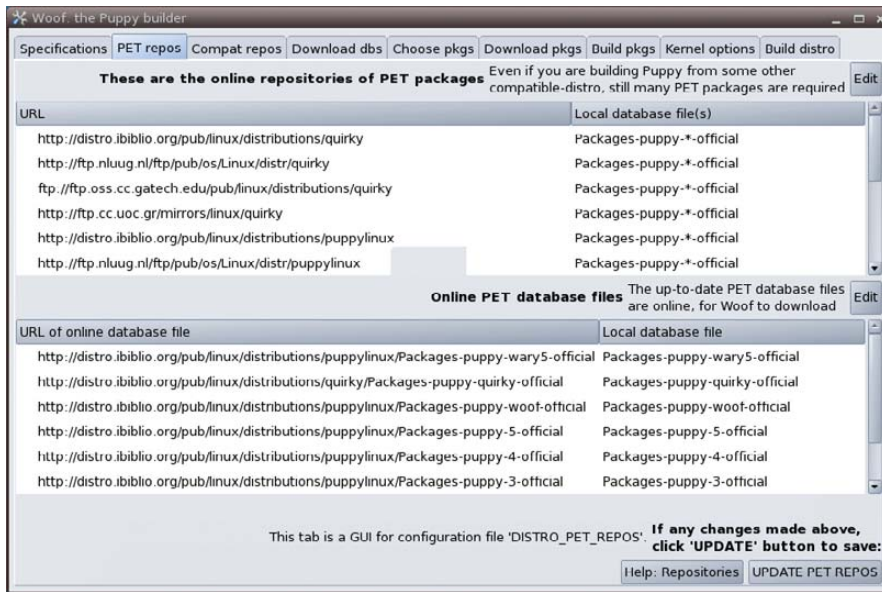




Figure 5. PET Repos

The tabs at the top of this utility correspond to the steps in building a distribution. The tabs mirror the Woof configuration files and allow you to

customize your distribution. To create your distribution, you move left to right through the tabs, making choices and executing the required builds.

The Specifications tab shows the high-level settings for the distribution to be built. At the time of this writing, woof comes preconfigured to build Wary, one of the derivatives of Puppy Linux. Notice the “Previous templates” section about two-thirds of the way down. From this pull-down, you can select a number of previously defined (and proven) build templates, which minimizes the choices you have to make and the possibility of specifying something that won’t work. I chose “Ubuntu Lucid Lynx LuPu”. If you make changes, select UPDATE OVERALL SPECS at the bottom to update the corresponding configuration files. You may get some errors about missing files—at this point, all of those should be configured for loading when Woof does the build.


The next tab, PET repos, tells Woof where to find the PET packages and databases it needs for the build. PET is the native package format used by Puppy Linux and Woof. It is optimized so the packages are small and quickly loadable. Unless you want to add a PET package that isn’t part of the standard Puppy PET



Develop.



Deploy.



Scale.

Full root access on your own virtual server for as little as \$19.95/mo

Multiple Linux distributions to choose from • Web-based deployment • Five geographically diverse data centers • Dedicated IP address • Premium bandwidth providers • 4 core SMP Xen instances • Out of band console access • Private back-end network for clustering • IP fail-over support for high availability • Easily upgrade or add additional Linodes • Free managed DNS

For more information visit www.linode.com or call us at 609-593-7103



linode.com

PET is the native package format used by Puppy Linux and Woof. It is optimized so the packages are small and quickly loadable.

repositories, you shouldn't have to change this tab. If you need to modify the list of repositories, click Edit, and you will be put into the text editor to edit the DISTRO_PET_REPOS configuration file. The Help: Repositories button will give you the format for the file. If you do make any changes, click UPDATE PET REPOS to save the updates.

The next tab, Compat repos, gives Woof the on-line locations for packages and databases compatible with the entry "Compatible-distro", that you made on the Specifications tab. If you need to modify the list of compatible package locations, click the appropriate Edit button, and you will be put into the editor to edit the DISTRO_COMPAT_REPOS configuration file. Again, the format is available with the Help button, and if you make changes, you need to click UPDATE COMPAT-DISTRO REPOS to save them.

The next tab, Download dbs, executes the script (Osetup) to download the PET and compatible package databases from the last two tabs. The top half of the tab shows the PET databases selected and those that are actually available locally. The bottom half does the same for compatible databases. If the selected and local lists do not agree (which is likely at this point), click on the UPDATE LOCAL DB

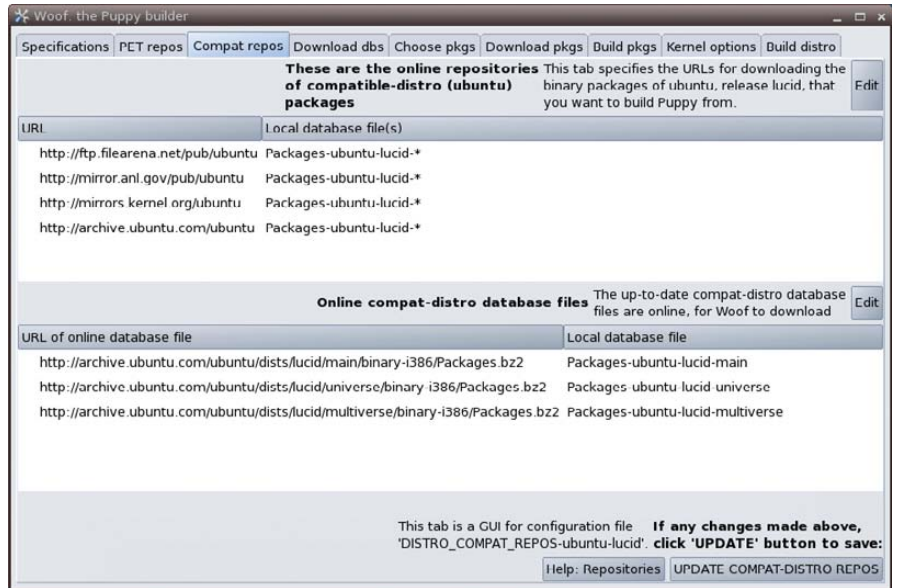


Figure 6. Compat Repos

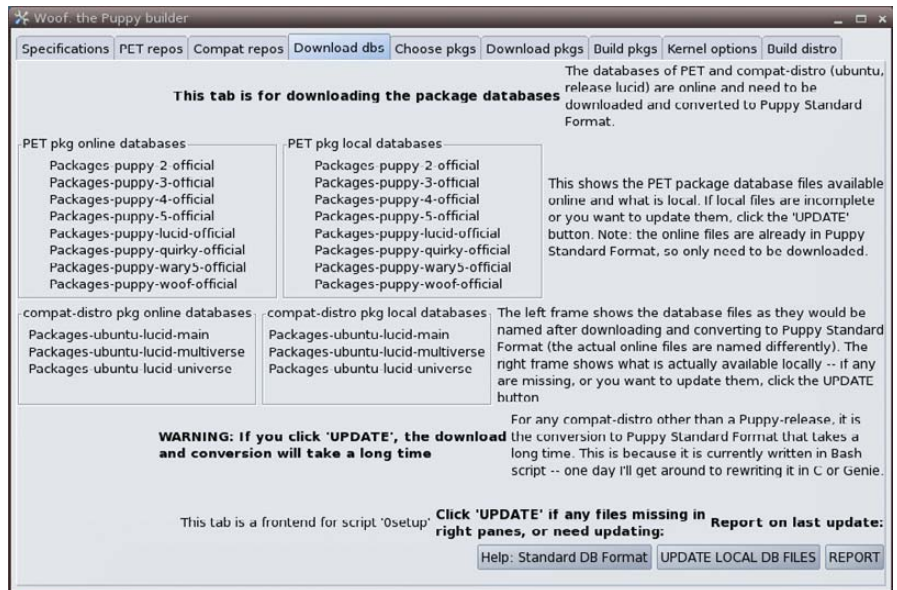


Figure 7. Download dbs

FILES button, and the files will be downloaded. This can take a while to complete.

The next tab, Choose pkgs, is where you see the PET and compatible packages included in your build. There are separate buttons to edit the list of PET packages and the list of compatible distro packages, but at the time of this writing, these editors have not been implemented. No matter, you can change the package selection by directly editing DISTRO_PKGS_SPECS-<compatible distro>. The format of the file is given as a comment in the file itself. At the time of this writing,

the CHECK DEPENDENCIES button also has not been implemented, so be sure the dependencies for any additions are included, so the build can complete successfully.

The Download pkgs tab executes the script (1download) that actually downloads the packages you specified for your distribution. Again, this takes a while. The script first checks to be sure the on-line package locations are reachable, then downloads each of the packages needed. Downloading can take several hours, so a small control panel dialog box pops up to let you pause, resume or quit the script. If

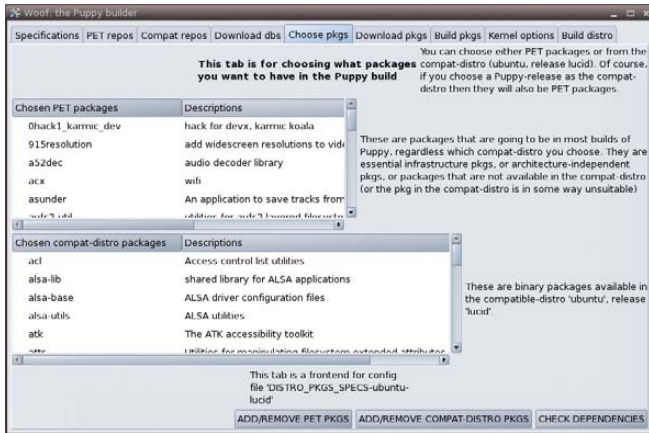


Figure 8. Choose pkgs

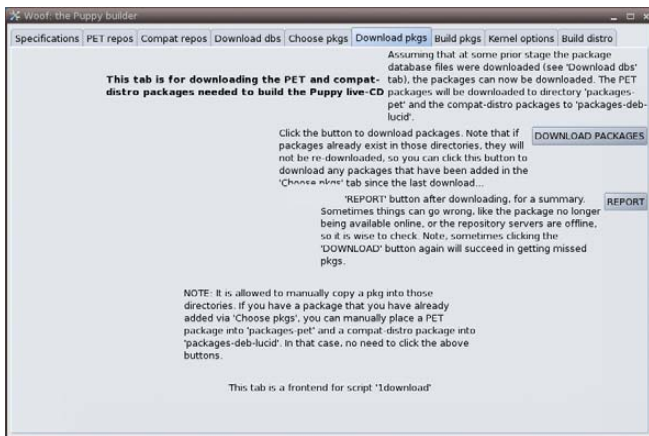


Figure 9. Download pkgs

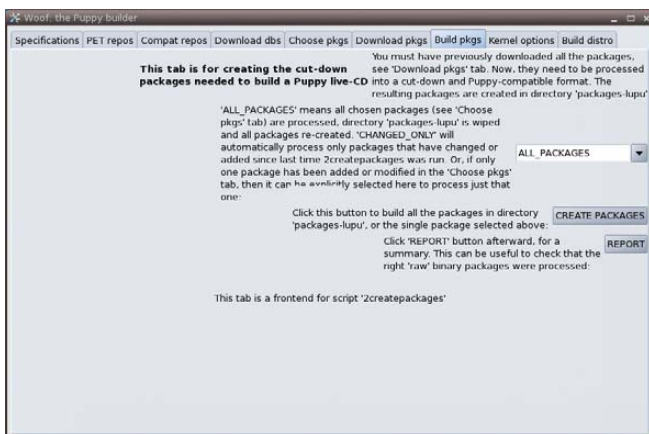


Figure 10. Build pkgs

a package already exists in the local location, it does not download again, so the script also can be used to restart a download or update the set of packages when a change is made. The REPORT feature on the tab is not yet implemented, but the results are in a file called DOWNLOAD-FAILS-PET or DOWNLOAD-FAILS-<distro> in the woof-tree directory.

Advertiser Index

CHECK OUT OUR BUYER'S GUIDE ON-LINE.

Go to www.linuxjournal.com/buyersguide where you can learn more about our advertisers or link directly to their Web sites.

Thank you as always for supporting our advertisers by buying their products!

Advertiser	Page #	Advertiser	Page #
1&1 INTERNET, INC. www.oneandone.com	C2, 1	LULLABOT www.lullabot.com	13
ABERDEEN, LLC www.aberdeeninc.com	3	MICROWAY, INC. www.microway.com	C4, 5
ANDROID DEVELOPER CONFERENCE www.AnDevCon.com	35	POSSCON www.posscon.org	67
DRUPALCON chicago2011.drupal.org	53	SCALE www.socallinuxexpo.org	17
EMAC, INC. www.emacinc.com	73	SILICON MECHANICS www.siliconmechanics.com	25, 57
GENSTOR SYSTEMS, INC. www.genstor.com	27	SOUTHWEST DRUPAL SUMMIT www.swdrupalsummit.com	9
IXSYSTEMS, INC. www.ixsystems.com	C3	SXSW FESTIVALS AND CONFERENCES www.sxsw.com	7
LINODE, LLC www.linode.com	63	TECHNOLOGIC SYSTEMS www.embeddedx86.com	15
LOGIC SUPPLY, INC. www.logicsupply.com	23		

ATTENTION ADVERTISERS

May 2011 Issue #205 Deadlines

Space Close: February 21; Material Close: March 1

Theme: Security

BONUS DISTRIBUTIONS: IT360, O'Reilly's MySQL, LinuxFest NW, USENIX NSDI

Contact Joseph Krack, +1-713-344-1956 ext. 118, joseph@linuxjournal.com

FEATURE Create Your Own Linux Desktop and Take It with You



Figure 11. Kernel Options

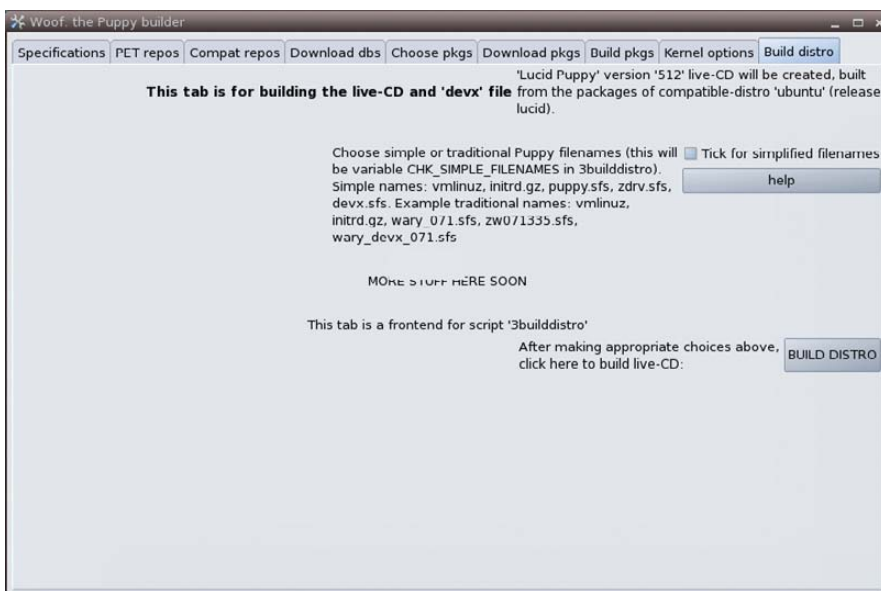


Figure 12. Build Distro

Here is where you get to use your Linux deductive skills if there were any problems in downloading the packages. In my case, the download for AbiWord failed—the script was looking for a pre-built PET package, `abiword-2.8.3-lupu.pet`. A quick `grep` of the `Packages*` files shows that the script expected to find it in `puppy-5-official`. Rerunning the `Download`, the problem is quickly apparent—`abiword-2.8.3` fails to download because of a broken link, and discussions on puppylinux.org confirm

that there are some issues with AbiWord packages. I could find the `.pet` or `.deb` packages and copy them in by hand, but to tell the truth, I don't need AbiWord in my distribution, so instead, I edited `DISTRO_PKGS_SPECS-ubuntu-lucid` and changed the `abiword` entry from "yes" to "no". `Download pkgs` confirms that everything is successfully downloaded.

The `Build pkgs` tab executes the script (`2createpackages`) that takes the distro packages you downloaded (Ubuntu `.deb` packages in this example)

and converts them to PET packages. At the end of this step, you have a complete collection of PET packages for your distribution.

Kernel options offers you a choice of Linux kernels and some options for paring down the kernel size. Depending on your concern about kernel size, your need for the specific modules mentioned and your willingness to experiment, you can choose whether to include them.

Stage 4: Build Your Distribution

`Build distro` is the final tab that uses all of the collected configuration and packages to build your custom distribution. You get a chance to choose simplified filenames on this tab, but reading the quite thorough Help file, there doesn't seem to be a solid reason for doing so. Clicking on `BUILD DISTRO` starts the build. As the build proceeds, you are given some options (moving modules to a separate file to improve boot time, default theme, text shadowing, Xorg drivers and executable stripping). You can experiment or take the default (the prompts will tell you if the default is not safe). When the `.iso` file is complete, you have the option of burning a CD—you'll need it, even to make a USB key. Then, you're asked if you want to build a `devx` file, which you'll need if you plan to do compiles and builds with your distribution. When the script is done, you'll find your distro's `.iso` (and the `.iso`'s MD5) in the directory `sandbox3`.

Stage 5: Install Your Distribution on a Bootable USB Key

Puppy Linux includes a USB drive installer that's very easy to use, but it knows how to install only the version that is running. That's why you needed to make a CD of your distro earlier. Boot that CD and plug in a USB drive. Click on `install` on the desktop. Choose `Universal Installer` from the first page.

Click `OK` twice. The dialog is self-explanatory, with options to correct things should a problem arise, and plenty of confirmation before actually writing to the USB drive.

Congratulations, you've just created your own Linux distribution, which you can hang on your keyring and boot on



Figure 13. Puppy Universal Installer

just about any PC you find. As long as the PC has at least 256MB and can boot from USB, you can boot your Linux desktop, do your work and not affect the underlying system.

So, there are several ways to get Linux on your keyring:

- Use the USB drive version of a major distro.
- Use a smaller, compatible distro, like Puppy Linux.
- Create your own unique distro using Woof.

With that many options, you can trade-off customization against effort to get exactly the right solution for your needs.■

Rick Rogers has been a professional embedded developer for more than 30 years. Now specializing in mobile application software, when Rick isn't writing software for a living, he's writing books and magazine articles like this one. He welcomes feedback on the article at portmobileapps@gmail.com.

4TH ANNUAL **COLUMBIA, SC**

PALMETTO OPEN SOURCE SOFTWARE CONFERENCE

PEOPLE FROM 14 STATES ATTENDED IN 2010.

2011 WILL BE BIGGER AND BETTER!

EARLY BIRD REGISTRATION AVAILABLE!

Columbia Metropolitan Convention Center

FOR MORE INFORMATION, GO TO
WWW.POSSCON.ORG
 OR CALL EVENT CHAIR AT (803) 240-1213

presented by
CONSORTIUM
 FOR ENTERPRISE
 SYSTEMS MANAGEMENT



POSSCON
PALMETTO OPEN SOURCE SOFTWARE CONFERENCE

MARCH 23-25, 2011
 at the
COLUMBIA CONVENTION CENTER
 Columbia, South Carolina

FEATURING:

LATEST OPEN SOURCE TECHNOLOGY

WORLD-CLASS SPEAKERS

HACKATHON'S AND BUG SQUASHES

RESUMÉ DROP AND "SUITS MEET SHORTS" EVENT

WORKSHOPS

FAMOUS FOR LOW REGISTRATION FEES AND WORLD-CLASS SPEAKERS

Kexi in Use

Kexi is an integrated data management application and forms part of the KOffice suite. It can be used for creating database schemas, inserting data, performing queries, processing data and producing reports. ADAM PIGG

This article takes new users through a range of concepts in Kexi, the KOffice database management program. The features available in Kexi range from simple table and query creation to more-advanced reporting and scripting. Kexi recently was released as part of KOffice 2.2. This article is aimed at new users and existing users of Kexi 1.6 as a demonstration of how useful Kexi can be.

Kexi 2 has taken three years of development to get to this stage from Kexi 1.6, and it missed the releases of KOffice 2.0 and 2.1 due to the small number of developers. Like many open-source programs, Kexi is developed solely in developers' free time as a hobby. It is hoped that this first release now is stable enough for use and that it will provide something to build upon for future releases.

To be a useful guide, it's helpful to work on a real use case. As the owner of the Kexi Facebook page (see Resources), I get weekly updates via e-mail with various statistics. The aim of this article is to get this data out of these e-mail messages and into a Kexi database to be able to perform queries and reports to show trends over time.

The Raw Data

To get the data for the database, I exported a selection of e-mail messages from KMail. This created a .mbox file, which basically is a text file containing all the messages. It would have been possible

Kexi file-based databases use SQLite as the underlying format, so they are readable by any SQLite-compatible program.

to go through each e-mail and enter the details into a table manually, but as I have a few built up already, I want to gather the data automatically. This provides a good challenge for writing a script within Kexi to do it for me.

Starting Off—Create a Database and Table

If Kexi is not included in your installation, see if it is available as an update in your package manager. If not, you need to install it from source using the guides on the KOffice and KDE Wikis (see Resources).

Start by launching Kexi and selecting create a Blank Database from the startup wizard. Depending on the installed plugins, you will be able to create a database stored as a file or create a database on an existing database server, such as PostgreSQL or

MySQL. Selecting to have it stored in a file is easiest for new users and is appropriate when there will be a limited number of users accessing the database at any one time. Kexi file-based databases use SQLite as the underlying format, so they are readable by any SQLite-compatible program.

The database requires a name (I chose kexi_facebook), followed by a location to save it. The default location is fine. After this, you are presented with the main Kexi window. The main window contains a toolbar along the top and a project navigator down the left-hand side. The main toolbar in Kexi is different from the other KOffice applications and uses a tab-style layout. Each opened window also has a local toolbar for options specific to that window, such as table, query, form, report and script.

From the Create tab across the top menu, choose Table to launch the table designer.

The statistics I receive via e-mail include the date, number of new fans, number of wall posts, number of visits and total fans, so I created a table with the design schema shown in Figure 1.

Field	Field Caption	Data Type	Comments
PK	id	Integer Number	
	stat_date	Date/Time	
	new_fans	Integer Number	
	new_posts	Integer Number	
	visits	Integer Number	
	total_fans	Integer Number	

Figure 1. Table Design

The fields have a name, type and comment, and also several properties are available from the property editor on the right-hand side, such as constraints and a default value if none is given. Each object in the database will have numerous properties associated with it, and the property editor allows these to be displayed and edited in a single location.

Switching to Data view prompts you to save the table and show the table data editor allowing manual entry of records, but that's not much fun!

Getting the Data

With my newly created but empty table, I needed to get the data automatically. As I mentioned earlier, the data was in a

single .mbox file containing all e-mail messages. Kexi supports scripts, which can be written in ECMAScript (aka JavaScript), Python or a number of other languages supported by Kross, the KDE scripting framework. I chose to use the QTScrip backend, which allows writing in JavaScript, as I am more familiar with it than Python.

My script had to open the .mbox file, read it line by line, grab the data it needed using string manipulation, and when a full set of data was read, add it as a record to the database. Scripts not only have access to built-in methods and Kexi-specific methods, but they also can import libraries containing large amounts of useful functions—the most useful being the Qt libraries. I use the Core functions to have access to the filesystem, using QTextStream for reading data, and the GUI functions for access to QMessageBox to present errors in a dialog if they occur.

From the Create menu tab, this time, I choose Script. This launches the script editor in the central window and the property editor down the right.

A script has only a few properties, the type and the interpreter. The interpreter I want is QTScrip, and the type is Executable. An executable script is one that is meant to be run manually. A Module script is meant to contain generic modules of code, accessible from other scripts. And, an Object script is one that is tied to another database object, such as a report.

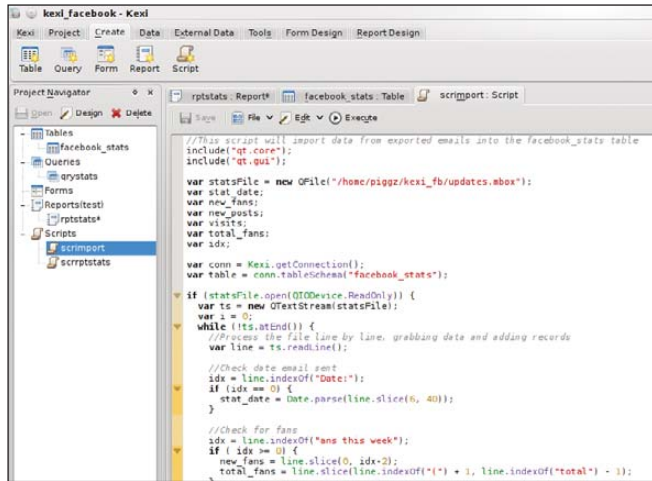


Figure 2. Script Design

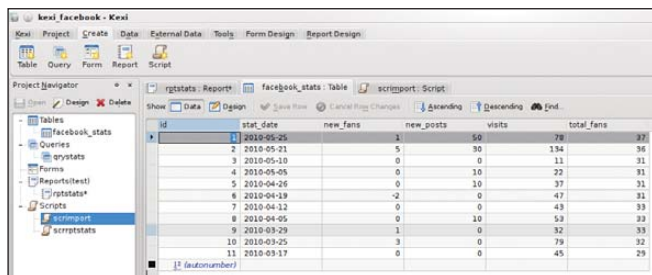


Figure 3. Table Data

The entire script is shown in Listing 1.

A possible bug in the script shown in Listing 1 is that it assumes there are no current records in the table, and it creates primary keys starting at 1. It is okay to run the script once, but if it is run again, it tries to overwrite records that have an ID matching what it is trying to insert. To make it more robust, it first needs to find out the current maximum of the ID field (this would be a good exercise to get used to writing scripts).

When executed from the script toolbar, the script gathered 11 records worth of data, which is visible from the Table Data View (Figure 3).

It's worth pointing out that the above script took a lot of trial and error, as it is not initially obvious that it is possible to import extra libraries or use Kexi-specific functions. The documentation

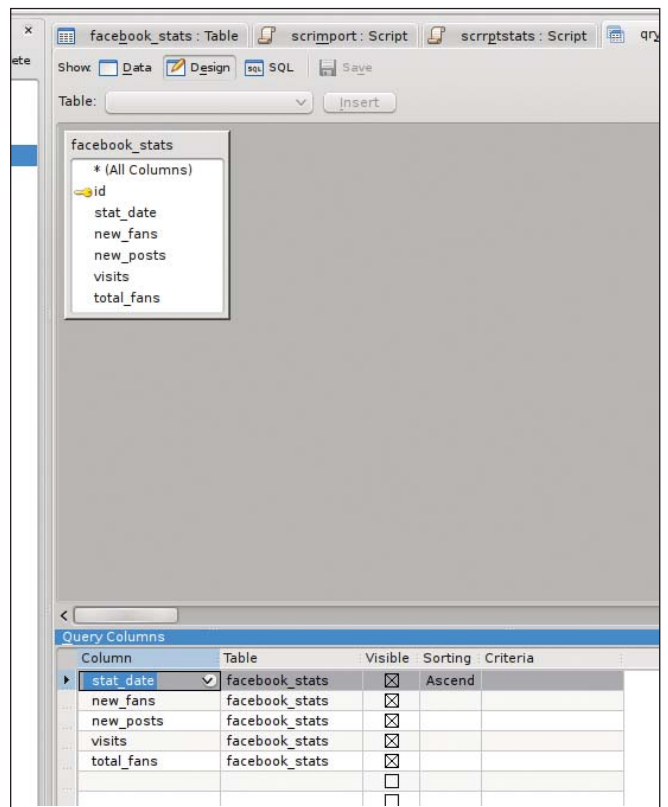


Figure 4. Query Design

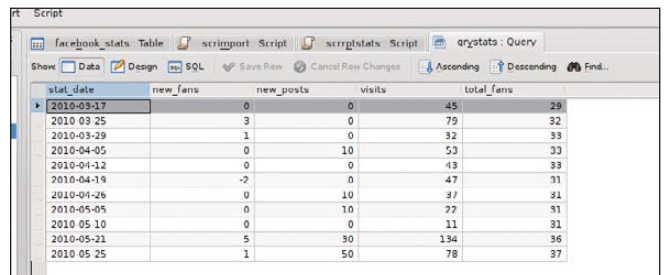


Figure 5. Query Data

Listing 1. Facebook E-mail Import Script

```
// This script will import data from exported emails
// into the facebook_stats table
include("qt.core");
include("qt.gui");

var statsFile = new QFile("/home/piggz/kexi_fb/updates.mbox");
var stat_date;
var new_fans;
var new_posts;
var visits;
var total_fans;
var idx;

var conn = Kexi.getConnection();
var table = conn.tableSchema("facebook_stats");

if (statsFile.open(QIODevice.ReadOnly)) {
    var ts = new QTextStream(statsFile);
    var i = 0;
    while (!ts.atEnd()) {
        // Process the file line by line,
        // grabbing data and adding records
        var line = ts.readLine();

        // Check date email sent
        idx = line.indexOf("Date:");
        if (idx == 0) {
            stat_date = Date.parse(line.slice(6, 40));
        }

        // Check for fans
        idx = line.indexOf("ans this week");
        if ( idx >= 0) {
            new_fans = line.slice(0, idx-2);
            total_fans = line.slice(line.indexOf("(") + 1,
                                   line.indexOf("total") - 1);
        }

        // Check for wall posts
        idx = line.indexOf("all posts");
        if (idx >= 0) {
            new_posts = line.slice(0, idx-2) + 0;
        }

        // Check for visits
        idx = line.indexOf("isits to your");
        if (idx >= 0) {
            visits = line.slice(0,idx-2);

            // Should have all the data now so insert a record

            stat_date = new Date(stat_date);
            var short_date = stat_date.getFullYear() + "-"
                + pad(stat_date.getMonth() + 1, 2) + "-"
                + pad(stat_date.getDate(), 2);

            if (!conn.insertRecord(table, [++i,
                                           short_date,
                                           new_fans,
                                           new_posts,
                                           visits,
                                           total_fans])) {
                var msg = "Cannot insert into " + table.caption() + '\n';
                msg += "Date: " + stat_date.toString()
                    + " " + short_date + '\n';
                msg += "New Fans: " + new_fans + '\n';
                msg += "Total Fans: " + total_fans + '\n';
                msg += "New Posts: " + new_posts + '\n';
                msg += "Visits: " + visits;

                QMessageBox.information(0,"Error", msg);
            }
        }
        QMessageBox.information(0, "Records Added:", i);
    }

    statsFile.close();

    function pad(number, length) {
        var str = '' + number;
        while (str.length < length) {
            str = '0' + str;
        }
        return str;
    }
}
```

needs work to make this easier for new users, and submissions are very welcome at the KDE Userbase Web site.

Sort the Data, Create a Query

At the moment, the data is ordered in the order in which it was extracted from KMail. Because I need it to be in ascending date order, I created a query to sort it. From the Create tab, this time I chose Query. I wanted all fields except the auto-incrementing

primary key, so I set it up as shown in Figure 4.

Switching to Data View executes the query and displays the results (Figure 5).

I saved the query as qryStats for use in a report.

Bringing It Together with a Report

A new feature of Kexi 2 is the report plugin. This allows reports to be designed and executed directly within Kexi

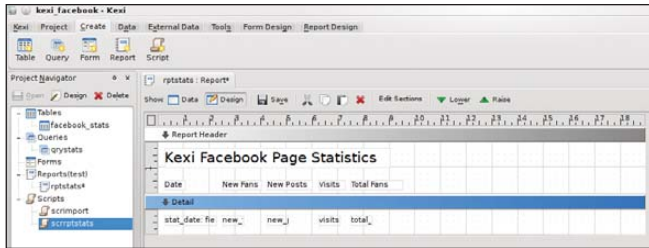


Figure 6. Report Design

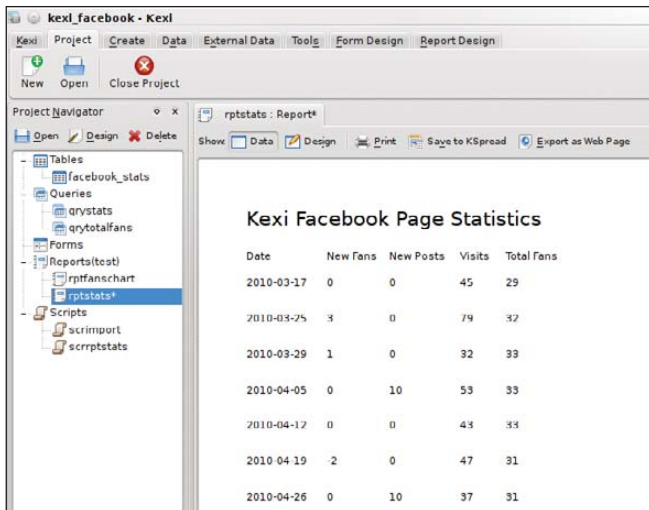


Figure 7. Plain Report

using a GUI editor similar to report designers in other database systems, such as Microsoft Access, Crystal Reports or Oracle Reports. In Kexi 1.6, reports were available as a separate add-on from kde-apps.org, but it did not contain as many features as the version in Kexi 2, and it was not fully integrated with the application, as the designer was an external program.

Reports can be printed, saved as a PDF, exported to HTML or OpenDocument Spreadsheet files or simply remain in the database for live viewing. It is possible to save the report in all these formats because of the two-stage generation process. Reports first are rendered into an intermediate description, and this description is used to generate the final version in whatever format is selected. In a future version, it is likely that extra formats will be supported, such as OpenDocument Text and XML, suitable for further processing using XSLT.

From the Create tab, I choose Report to create a blank report with a single "Detail" section. The structure of a report is based around Sections, which can be page headers or footers, report header or footer, or Group sections where data is grouped on a field value.

Initially, all I want is a simple tabular view of the data, so all the fields will go into the detail section, apart from a header, and the field titles, which must go either in a Page Header or Report Header. From the Section Editor on the report toolbar, I added a Report Header, and using the Report

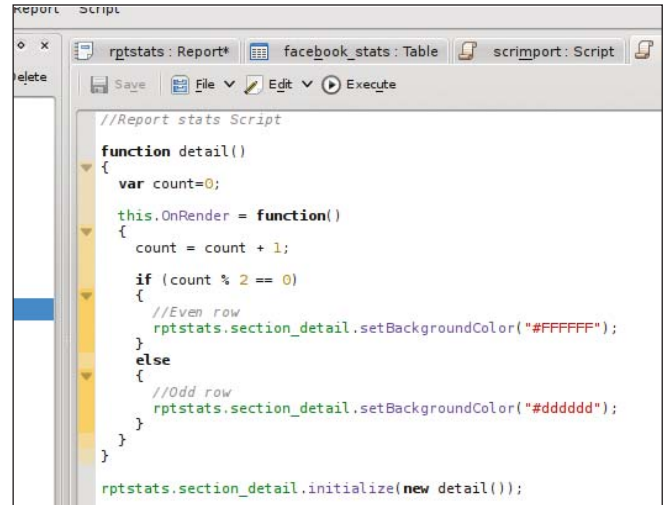


Figure 8. Statistics Script

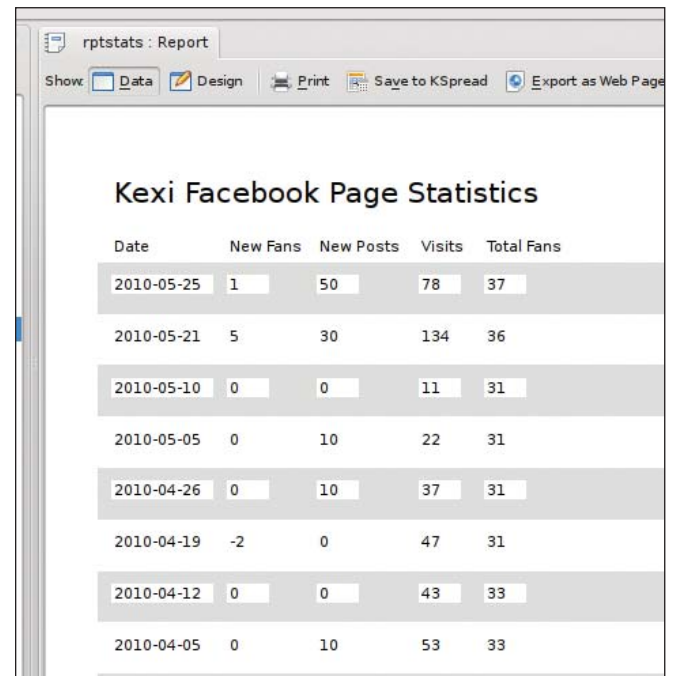


Figure 9. Tabular Report 1

Design tab on the menu bar, I added fields and labels to create the report layout. From the Data Source tab on the sidebar, I set the reports data source to the qryStats query I created above. Finally, I set the Control Source property of each field item to the corresponding field in the query and the Caption of the labels appropriately. In the end, it looked like Figure 6, and it generated a report, shown in Figure 7.

This gets the job done, but it isn't quite as "jazzed up" as I would like. It's common in desktop applications to alternate the background color of rows to make it more obvious where each set of data begins and ends, so let's try that.

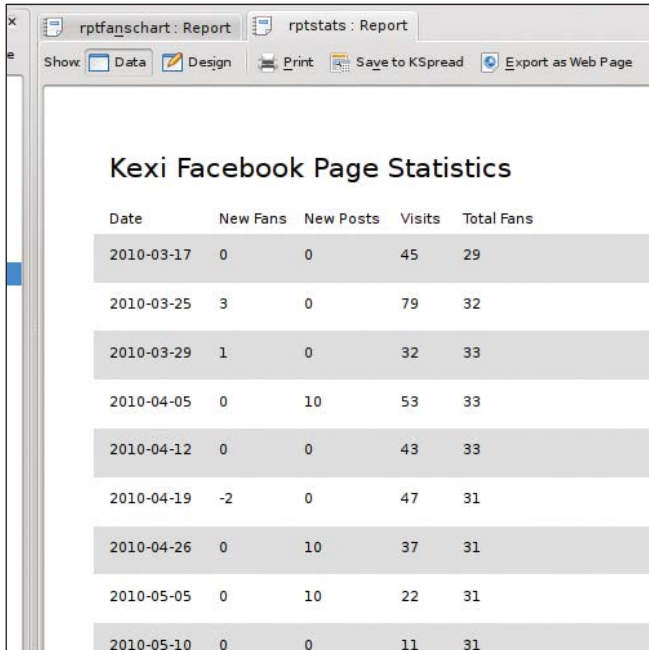


Figure 10. Tabular Report 2

Report scripts use the features of Kross::Object, where each object in a report can be associated with a script object, making it more object-oriented in nature.

I created another script, but this time set its type to Object, as it is to be associated with the report object. Report scripts are event-driven—that is, whenever a certain event occurs in the generation of the report, the associated code in the script is called. Report scripts use the features of Kross::Object, where each object in a report can be associated with a script object, making it more object-oriented in nature. Each script object can have its own variables and functions. Report objects can be the report itself or any of the report sections. To make it more clear, the final script looks like what's shown in Figure 8.

This is quite a simple script. There is an object called detail, containing a function OnRender, which will be called whenever a detail section is rendered. The function keeps track of how many times it has been called and alternates the background color. The final line of the script associates the detail function with the detail section of the report.

Then, in the report, I set the Interpreter Type to QTScript and the Object Script property to the name of the script. It is important that the Interpreter type of both the report and script match; otherwise, the script won't be presented as an

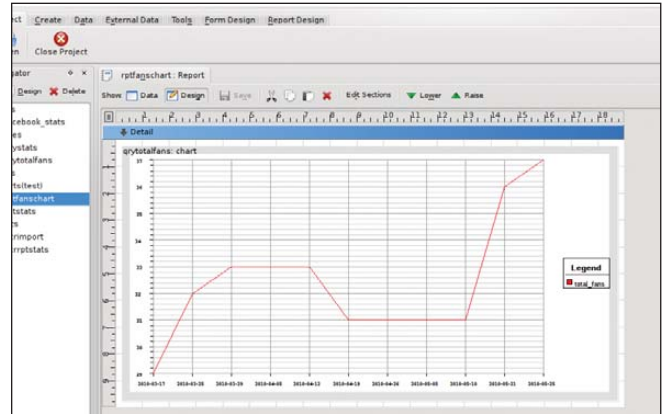


Figure 11. Chart Design

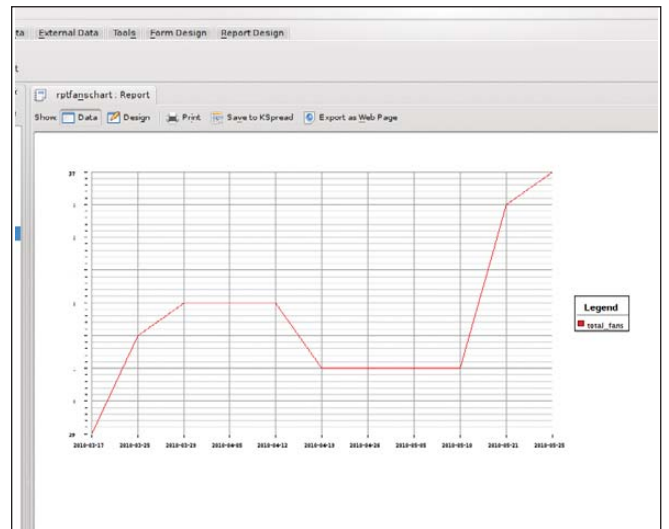


Figure 12. Chart Report

option in the Object Script list.

The generated report now looks like Figure 9.

It's not so great with the white background on the fields, so I went back to the designer and changed the Opacity property of each of the fields to 0 to make them transparent, resulting in a more reasonable report (Figure 10).

Adding Something Trendy

My final requirement at this stage was to have something more graphical: a nice chart to show the trend of fans over time. The report designer allows the creation of charts using the KDChart library from KDAB and is used in the KOffice program KChart. It is quite powerful, allowing you to join chart data to the main report data (called master-child links), but for now, all I needed was a simple, single chart. The chart object expects data in a certain format. There must be two or more columns of data. The first column is used for the labels on the X axis, and all other columns are used as a series

	A	B	C	D	E	F	G
1	Date	Kexi Facebook	New Fans	New Posts	Visits	Total Fans	
2	2010-03-17	0	0	45	29		
3	2010-03-25	3	0	79	32		
4	2010-03-29	1	0	32	33		
5	2010-04-05	0	10	53	33		
6	2010-04-12	0	0	43	33		
7	2010-04-19	-2	0	47	31		
8	2010-04-26	0	10	37	31		
9	2010-05-05	0	10	22	31		
10	2010-05-10	0	0	11	31		
11	2010-05-21	5	30	134	36		
12	2010-05-25	1	50	78	37		
13							

Figure 13. KSpread Data

in the chart. I started by creating a query with two columns, date in ascending order and total fans, then created a new report. The report itself is not based on any data, so its Data Source was left empty. An empty data source will produce a report with one detail section, providing an area to add a minimal set of items to a report.

In my detail section, I added a chart object from the report designer toolbar and set its data source to the query I had just produced (Figure 11).

As you can see, even at design time, the chart object is able to gather data and draw a preview of the chart. Switching to the data view shows the chart without any of the extra lines and text from the designer (Figure 12).

Hard Copies

When printed, both the tabular report and chart report look as they do in the Data view. When printed using the PDF printer option in KDE, the chart even retains all its detail, as it is not converted to a bitmap, but saved as lines, which makes it completely zoomable!

Saving the tabular report as an HTML document produces two options: saving as a table or using CSS. The table option produces an HTML file where the text from each field in a report is saved as a table cell, and each section is a row. The

CSS option uses the <div> tag and tries to create an HTML file that closely resembles the original, allowing text and images to be rendered at arbitrary positions.

The tabular report also exports nicely into an OpenDocument Spreadsheet file for use in either KSpread or OpenOffice.org (Figure 13).

As you can see from the image, one problem is that the report's title has taken a cell with the other field headings. This is because it is in the same section, but it easily can be fixed by putting the title into a separate section, such as a Page Header.

Kexi 2 is a powerful tool for developing database applications. Find out what else is possible by giving it a try, and if you can, please contribute more documentation at userbase.kde.org/Kexi, or join the team by dropping into #kexi or #koffice on Freenode IRC. ■

Adam Pigg is a software engineer who, in his free time, contributes to the Kexi and KOffice projects. He's been contributing to KDE in general for around eight years. His remaining time is occupied by his wife, his four great kids and mountain biking. He usually can be found in #kexi on Freenode, where he will do his best to answer any queries.

Ultra Small Panel PC

PPC-E4

- Fanless ARM9 200MHz CPU
- 3 Serial Ports & SPI
- Open Frame Design
- 2 USB 2.0 Host Ports
- 10/100 BaseT Ethernet
- Audio Beeper
- Micro SD Flash Card Interface
- Battery Backed Real Time Clock
- 64 MB Flash & 64 MB RAM
- Linux with Eclipse IDE or WinCE 6.0
- JTAG for Debugging with Real-Time Trace
- WQVGA (480 x 272) Resolution TFT LCD with Touch Screen
- Four 12-Bit A/Ds, Two 16-Bit & One 32-Bit Timer/Counters

The PPC-E4, an ultra compact Panel PC with a 4.3 inch WQVGA (480 x 272) TFT color LCD and a resistive touch screen. The dimensions of the PPC-E4 are 4.8" by 3.0", about the same dimensions as that of popular touch cell phones. The PPC-E4 is small enough to fit in a 2U rack enclosure. **Price is \$345 at quantity 1.**

For more info visit: www.emacinc.com/panel_pc/ppc_e4.htm

Since 1985
OVER
24
YEARS OF
SINGLE BOARD
SOLUTIONS

EMAC, inc.

EQUIPMENT MONITOR AND CONTROL

Phone: (618) 529-4525 • Fax: (618) 457-0110 • www.emacinc.com

Resources

Kexi Project Facebook Page: www.facebook.com/kexi.project

Building KOffice: wiki.koffice.org/index.php?title=Building/Building_KOffice

Getting Started/Build/KDE4: techbase.kde.org/Getting_Started/Build/KDE4

Social from the Start

The Ayatana Project is striving to improve the perception and presentation of information on the desktop. The fruits of the project can be seen in the way that social media and social networking are increasingly integrated into the core of the Ubuntu desktop. JONO BACON

From the beginning of the evolution of Linux on the desktop, one of the many benefits of the open-source methodology is that we often can be reactive to user needs much quicker than our proprietary counterparts. The reason is simple and gets to the heart of open source itself; with more people inspired, motivated and equipped to solve problems, problems are solved more quickly, and everyone in the community benefits.

Just over a year ago in Spain at the Ubuntu Developer Summit, Mark Shuttleworth, founder of the Ubuntu Project, started coining the term Social from the Start. His idea was also simple: to build social media and social networking into the core of the Ubuntu desktop. As the leader of the Ayatana Project (which has been driving desktop innovation in Ubuntu), part of his vision is that access to social tools should be simple, elegant and integrated. We have seen the true fruits of these efforts shipped in the Ubuntu 10.04 Lucid Lynx release.

In this article, I peel back the covers and explore what Social from the Start currently includes in the Ubuntu 10.04 Lucid Lynx desktop and the opportunities we have open to us in the future.

The Messaging Menu

When the Ayatana Project first was announced, its focus clearly was articulated as improving the perception and presentation of information in the desktop, hence the name of the project—the Buddhist term for a “sense base” or “sense sphere”. The first innovation along this mission was the new notification bubbles that have shipped with Ubuntu for a few releases now. The next major change was the messaging menu, as shown in Figure 1.

The messaging menu provides a single place in which new information is made available to you. In a default Ubuntu installation, e-mail in Evolution, chat messages in Empathy and tweets in Gwibber are all made available in the messaging menu. In the past year, we also have seen numerous other applications make use of the messaging, such as the XChat-GNOME IRC client and Zimbra. The messaging menu neatly merges all of these different information flows together into one common place on your panel, only ever a click away.

By default, the messaging menu includes three primary types of content:

- Chat: interactive text-based real-time discussions with friends and colleagues.
- Mail: e-mail messages.
- Broadcast: social-networking broadcast messages, such as Twitter.

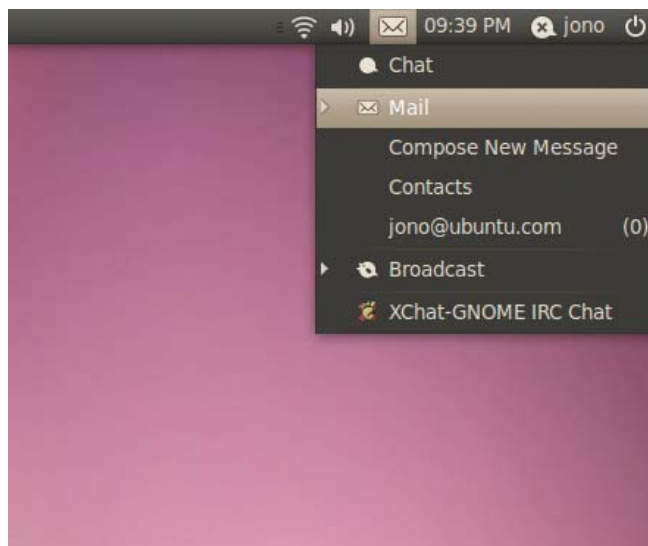


Figure 1. The Messaging Menu's Notification Indicators

Broadcast

A key component at the core of the social-networking support in Ubuntu is a simple little tool called Gwibber. Gwibber is a window in which to view a world filled with different social-networking Web sites and streams. Gwibber provides support for Facebook, Twitter, identi.ca, Flickr, StatusNet, FriendFeed, Digg, Qaiku and more networks are added with each release.

To use Gwibber, you first need to have an account on at least one social-networking service it supports. As an example, I have a Twitter feed at www.twitter.com/jonobacon, and I can use Gwibber to send and receive content without ever having to use the normal Twitter Web site. Gwibber's major benefit is that it brings all of these different social-networking services together into one window. Continuing my example, in addition to Twitter, I also have a Facebook page (www.facebook.com/jonobacon) and an identi.ca account (identi.ca/jonobacon). Ideally, not only do I want to read all of these feeds in one place, but also when I post a message, I want it to appear on all three at the same time. This is simple with Gwibber.

Let's first set up your accounts for the system. To do this, click your user name in the desktop panel (in the top-right part of the screen), and in the menu, select Broadcast Accounts.... The window shown in Figure 2 appears.

Click the Add... button, select the type of account and then



Figure 2. Setting up broadcasts accounts is simple in Gwibber.

click the Add button next to it. Finally, enter the login credentials for the account, and click the Add button to complete. Repeat this process for each of the different accounts you want Gwibber to talk to.

Now, let's see your accounts in action. Click the messaging menu (the little envelope in the panel), and click Broadcast to load Gwibber. A window that looks similar to Figure 3 appears.

Gwibber is split into a series of panes called streams. The stream on the far left is the Messages stream. It displays all the messages of the people you follow across the accounts that you set up in Gwibber. It merges all of these different messages into one place and organizes them chronologically. You can see which network a message is from by looking at the icon next to the name of the person in the message. Gwibber repeatedly updates all the different streams to keep you up to date with new content.

To the right of the Messages stream is the Replies stream. These are all the messages directed at you (for example, in my Gwibber, all messages with @jonobacon appear in this stream). Messages that come into this stream not only appear in Gwibber but also a notification bubble appears indicating a message has arrived that you may want to reply to.

To reply, hover your mouse over the message in Gwibber, and click the small envelope. The recipient's user name now appears in the text entry at the bottom of the window. Type your message, and Gwibber indicates how many letters you have left within the 140-character limit. With the character limit so precious, Gwibber also handily shortens Web addresses you paste into it. When you are ready, press Enter or click the Send button to send your message. Your message appears in your Messages stream to show it was posted.

When sending a new message, under the message entry text box is an icon next to Send with for each of the different accounts you configured

in Gwibber. Click these icons to select or deselect to which networks your message will go. By default, your message will go to all networks.

Don't Cross the Streams

One of Gwibber's most useful features is its support for multiple streams. With it, you can provide quick access to different themes of messages you want to see. This most typically includes searches for all messages, including particular search terms. It's incredibly handy if you want to see what the microblogging world is saying about something you care about, such as your name, product, project or service.

Setting up a new stream is simple. In Gwibber, click Gwibber→New Stream, and a new pane appears to the right of your existing streams (you may need to move the scroll bar to see the new view). In the new stream to the left is a down arrow, click it and a window pops up where you can select what will go in the stream. At the top are general categories of information common to all networks, such as Messages, Replies and Images. Underneath them are each of the different networks and the different types of information specific to those networks. This makes it simple to have a stream that shows only your Twitter replies, for example.

At the bottom of the list of content types is a Search option. Clicking takes you back to the main Gwibber window, and a text box appears where you can enter a search term. Add something and click the Search button, and a new stream to the right appears showing all messages across your networks with that search in them.

Tweeting from Your Panel

One of the key enhancements that Ubuntu 10.04 Lucid Lynx brought was the ability to tweet directly from your desktop without even needing to access the Gwibber window. With your broadcast accounts all set up, you can tweet by clicking your user

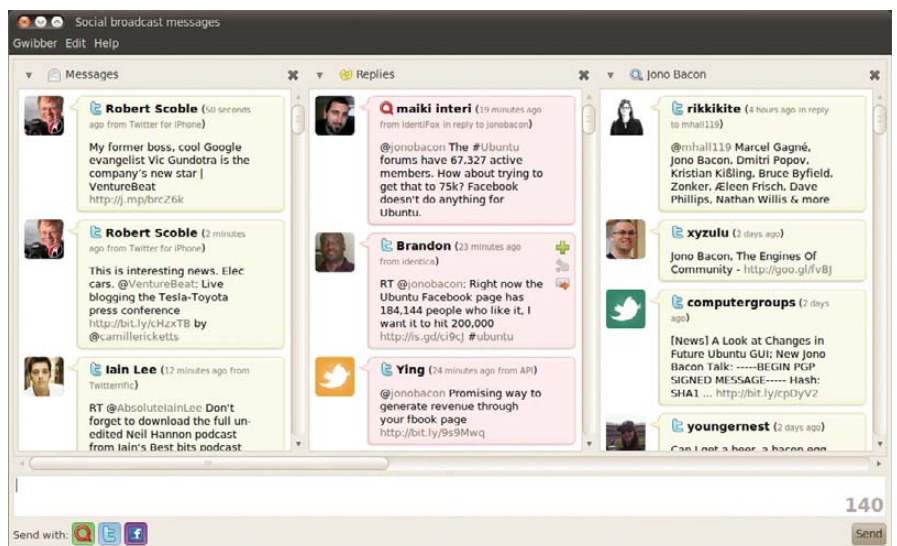


Figure 3. Gwibber provides a stream of social-networking messages.



Figure 4. Send tweets right from your desktop!

name (you also can press Super-S) on the panel and typing your tweet into the text entry box in the menu (Figure 4). This makes tweeting from Ubuntu a snap; one click, and your thoughts are broadcast instantly!

Chat

Another key component in building social features into the Ubuntu desktop is Empathy—a multiprotocol, extensible and powerful chat client. With Empathy, you can chat to your friends across a range of different networks all in the same place. Empathy supports Facebook Chat, Google Talk, Jabber, AIM, gadugadu, GroupWise, ICQ, IRC, MSN, mxit, MySpace, qq, same-time, silc, SIP, Yahoo!, Yahoo! Japan, zephyr and people who are nearby to you on your local network. Empathy brings all of these different networks together into one window, providing a single consistent user interface for all of your friends.

To get started with Empathy, first you need to set up your various chat accounts. Click your user name in the panel and then click Chat Accounts. A window that looks very similar to the Gwibber account setup window appears. This interface is almost identical; click the Add button, select an account type and add your account details. With each of your accounts set up, you now can fire up Empathy by clicking the messaging menu and clicking Chat.

Empathy provides a simple interface for accessing all your chat contacts in one place (Figure 5). For all contacts in your list, the icon on the left side indicates their availability. Green means they are available, orange means they are away from their keyboard,

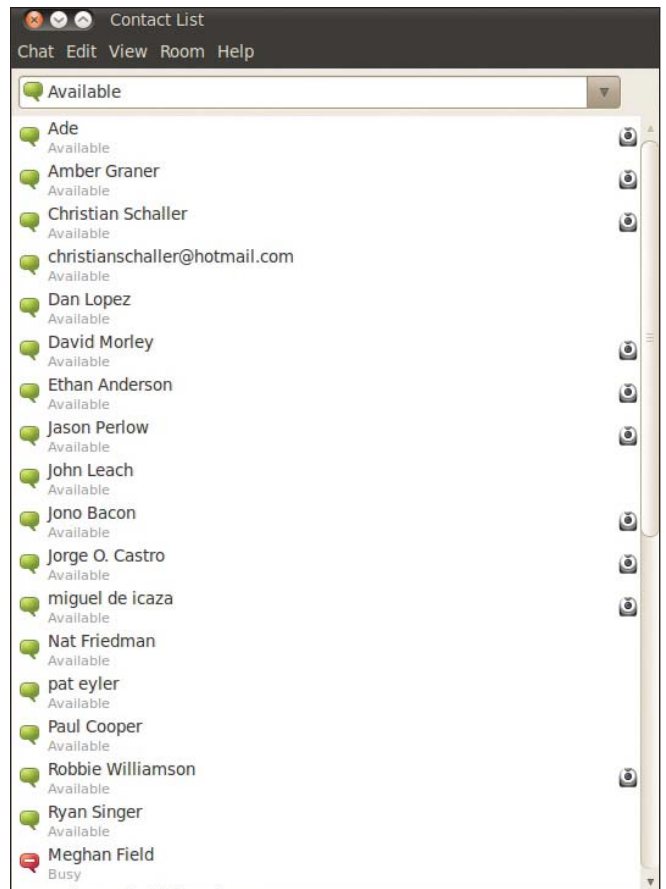


Figure 5. Empathy brings together all your chat contacts into one window.

and red means they are busy. To add a contact to your list, click Chat→Add Contact, select the appropriate network, and then add their credentials. They appear on your contact list when they approve your request.

To send an instant message to one of your contacts, double-click his or her name in Empathy and a new window appears where you can chat. Within the chat window, you should know about a few useful functions. First, if you want to find out a little more information about the contact, click Contact→Information. Next, you may want to send that person a file, which you can do by clicking Contact→Send file. Finally, if you ever want to look at previous conversations with your contacts, click Contact→Previous Conversations. This opens up a new window with a calendar and a list of contacts, and you can click a contact and skip through the calendar to browse logs of your discussions.

Earlier, I discussed how you can see someone's availability, so I also should explain how to set yours. There are two ways to indicate if you are available, away, busy or invisible (this is when you can see your contacts, but they can't see you). First, in the Empathy contacts window, use the drop-down box at the top of the window. Second, you can go to the Me Menu (the panel menu with your user name) and select one of the options there.

Audio, Video and Desktop Sharing

Although on the surface Empathy may seem a powerful and flexible instant-messaging client, but limited to textual messaging, it also packs in some other interesting features. The Ubuntu 9.10 Karmic Koala shipped Empathy for the first time with its audio and video support. With it, you can have live voice and video calls with your contacts.

Audio chat works with most networks and people who are connected to Empathy with a working microphone. For those contacts with whom you can videochat, they are indicated with a small Webcam icon next to their names in the Empathy contact list. To start a call, simply right-click a contact and select either Audio Call or Video Call. A new window appears and the call is initiated.

Another feature that a lot of people are unaware of in Empathy is desktop sharing. With it, you can share your desktop with contacts so they can control it remotely, moving your mouse and typing on your desktop. This is handy for helping people fix problems or for collaborating on projects. Sharing your desktop requires a fairly meaty Net connection on both sides, so don't try to use it if you have limited bandwidth. If you want to give it a try, simply right-click a contact and select Share My Desktop.

Wrapping Up

With Ubuntu 10.04 Lucid Lynx, the Social from the Start initiative has really started to embed into the desktop and operating system. With it, you can tweet from your desktop to a wide variety of networks, have text, audio, video and desktop sharing support with all of your contacts in one place, and aggregate the many different messages that you need all in one place—the messaging menu. These different components have been designed to fit neatly together, sharing many of the design and interaction characteristics to provide a smooth, consistent and fluid user experience.

Although the experience is exciting in itself, the machinery under the hood is even more compelling. Empathy and Gwibber are at the forefront of innovation in messaging. Empathy, part of the GNOME desktop, is based on the Telepathy framework, a powerful set of tools for communicating over different networks and mediums, and aggregating these different mediums together. Although Empathy is the front end, the underlying Telepathy framework is opening up tremendous opportunities for applications in the future. We can thank Collabora Multimedia in England for much of the work on this.

In the same vein, at the heart of Gwibber is the Gwibber API. Although Gwibber provides the user experience we care about, the Gwibber API fulfills a means for application developers to tweet from their application. I myself used this in a program, and I added support to tweet a message with merely three lines of Python. These technologies provide developers with the ability to build modern social-networking features into their applications, continuing to build huge value in the open-source desktop.

It is an exciting time in the Linux desktop world, and this article has covered merely the beginning of the journey. This may be the start, but there is a long road ahead of us—it is time to get ever more social. ■

Jono Bacon is the Ubuntu Community Manager at Canonical, author of *The Art Of Community* published by O'Reilly, founder of the Community Leadership Summit and co-presenter on *Shot Of Jaz* and *FLOSSWeekly*.

**NOW ON
SALE:
SAVE 15%***

The newly updated
LINUX JOURNAL ARCHIVE
is here!



The archive includes **all 200 issues of Linux Journal**, from the premiere issue in March 1994 through December 2010. In easy-to-use HTML format, the fully searchable, space-saving archive offers immediate access to an essential resource for the Linux enthusiast: *Linux Journal*.

*Use coupon code **ACDNEW** at www.linuxjournalstore.com, now through March 31, 2011, and save 15% off the standard retail price of \$34.95.

www.LinuxJournalStore.com



KYLE RANKIN

Tablets

Who wants one?

BILL: Hey, Kyle, what do you think of tablets?

KYLE: Well, I think they are fine for storing commandments, but for computing, I think they are a bit limiting. I take it you probably think they are great.

BILL: I don't think they're great, but they certainly have come a long way. In the right setting, I can see how they'd be very useful. I'm thinking about getting one myself.

KYLE: See, when I use a computer, I tend to type a lot. I'm pretty particular about getting tactile feedback in my keyboards whether it's for a desktop (Model M thank you) or a laptop. Typing on a flat surface with no feedback just seems like a step backward.

BILL: I remember, you're all about the battleship keyboard. I know that having your keyboard survive a zombie insurrection is a major selling point for you. However, there's more to a tablet than just typing.

A laptop or Netbook is the tool of choice for content creation, but for content consumption, I think a tablet might be the way to go.

Having an instant-on Web browser and mail client just kicking around would be very handy.

KYLE: You mean like a laptop that can suspend to RAM? Or a Netbook?

BILL: I have yet to see any laptop that's truly instant-on, even when suspended to RAM. They always take a few seconds to get on their feet. And even then, the form factor is different. A laptop or Netbook is the tool of choice for content creation, but for content consumption, I think a tablet might be the way to go. I know more often than not I'll reach for my iPhone if I need to check mail real quick. It's just faster and more convenient.

KYLE: See, that's just the point, I think the tablet has long been a solution in search of a problem. Now it has to compete with a smartphone for portable, underpowered computing, a Netbook for inexpensive portable computing, and a laptop or desktop for full-featured computing. Having a large fruit-named company create one (and new companies throw cell-phone software on their

tablets) doesn't change that. I had a hybrid laptop that could rotate into tablet mode, and I think I used it maybe a handful of times, and even then, it was just as a novelty e-book reader.

BILL: I remember, but that laptop wasn't exactly a powerhouse either. And if I recall, it was running stock Ubuntu, which is not a portable-optimized OS like Android or iOS. I think your definition of computing is different from that of a lot of folks, Kyle. As a system administrator and writer, your use case depends on having multiple windows, a full-size keyboard, and the storage and horsepower of a conventional laptop. However, as iPad sales prove, there's a huge segment of the population who just wants to surf using tablets and play *Angry Birds*.

KYLE: Unless you are wearing some interesting jeans, a tablet isn't going to be any more portable than any other similarly underpowered Netbook, but you'll pay a premium for the fingerprint-smear

touchscreen and the lack of a keyboard. I think even surfing suffers on a tablet. However hyperlinked the Web might be, these days, people keep talking about everyone "contributing to the conversation" and other Web 2.0 terms. It's hard to do that just by touching and dragging on a screen.

BILL: Actually, SCOTTEVEST makes a vest that can hold an iPad in an internal pocket like a holster, but that's a little much even for me. Considering how I'd use a tablet, I wouldn't miss the keyboard. It wouldn't be my primary computing device, nor would it be what I reached for if I needed to do heavy work.

KYLE: And no matter how sophisticated your touchscreen keyboard is, it's still a keyboard on a flat surface with no tactile feedback. Plus, you lose a good portion of that tablet screen for it. I'd still love to see a picture of all of your active computing gear stacked on top of each other in size order. I think your iPad would fit somewhere on top of your small laptop but underneath your Netbook, Nokia pocket tablets and smartphones.

LINUXTM JOURNAL

Linux News and Headlines Delivered To You



Linux Journal
topical RSS feeds
AVAILABLE

BILL: I think you can do more on a tablet than you realize. I've taken two trips with nothing but my iPhone and I actually *survived* just fine. A tablet would be nice to take on those trips for the great battery life and larger screen. For me, I think it fits between a laptop and iPhone nicely. You keep asking for pics of my computers. It's like you're just interested in me for my system porn. I HAVE A MIND TOO.

KYLE: I've taken a number of trips myself with just an N900, and the only things I've missed were a full-size keyboard and larger screen. Anything much larger than that, and you lose the real benefit of a portable smartphone. Basically, to me, anywhere you can take a tablet, you can take either a Netbook or a laptop. Although touchscreen interfaces are all the rage now, I still think they are really limited compared to a lot of the traditional input methods. In my mind, if you truly do want a tablet, the smart move is to invest in a convertible model, so you at least can get a proper hardware keyboard when you need one.

BILL: That's where a tablet fits in. I bring my Nook now on trips as a book, and it works great. If that device had a faster CPU, faster screen refresh and the ability to do a little more, I would carry it around constantly. I don't need it to do everything; I just need it to do 90% of everything. You can pair a Bluetooth keyboard to a tablet, and then you've got your proper hardware keyboard.

KYLE: For three hours until the battery dies, at least....

BILL: Three hours? Now you're resorting to hyperbole. The Galaxy Tab is good for seven, and the iPad is good for ten, according to specs and reviews I've seen. I know my Bluetooth keyboard batteries last a lot longer than either of those.

KYLE: The bottom line for me is that for completely portable (and touchscreen) computing, I have an N900 with a physical keyboard. For full-featured computing, I have a rather portable laptop. To me, and I imagine to most of the population, a tablet is nothing more than the third computer no one really needs.

BILL: I think the market proves your opinion to be wrong, Kyle. Apple sold 4.2 million iPads in the last quarter—4.2 *million*. There's a market for tablets and a use case for them. It may not be something *you* want, but the public's definitely interested, and so am I. I'm going to see what the next generation of iPad and Android tablets brings to the table before I pull the trigger on a purchase, but I definitely can see myself using one for light-duty computing. ■

Kyle Rankin is a Systems Architect in the San Francisco Bay Area and the author of a number of books, including *The Official Ubuntu Server Book*, *Knoppix Hacks* and *Ubuntu Hacks*. He is currently the president of the North Bay Linux Users' Group.

Bill Childers is an IT Manager in Silicon Valley, where he lives with his wife and two children. He enjoys Linux far too much, and he probably should get more sun from time to time. In his spare time, he does work with the Gilroy Garlic Festival, but he does not smell like garlic.

http://www.linuxjournal.com/rss_feeds



Dysintegration

The open market is as big as its many fragments. DOC SEARLS

During Apple's Q4 2010 earnings call with analysts, Steve Jobs did a typically masterful job of re-clothing an old argument in fresh duds. The old argument was open vs. closed. The new one was integrated vs. fragmented. Time will tell if the clothes will fit. Meanwhile, I'd like to devote this space to helping make sure they don't. So, here's what Steve said, with corrective replies. From a transcript of the call, in tech.fortune.com (tech.fortune.cnn.com/2010/10/18/what-steve-jobs-said-about-google):

Google loves to characterize Android as open, and iOS and iPhone as closed, we find this a bit disingenuous and clouding the real difference between our two approaches. The first thing most of us think about when we hear the word open is Windows, which is available on a variety of devices. Unlike Windows, however, where most PCs have the same user interface and run the same app, Android is very fragmented. Many Android OEMs, including the two largest, HTC and Motorola, install proprietary user interfaces to differentiate themselves from the commodity Android experience. The users will have to figure it all out. Compare this with iPhone, where every handset works the same.

Calling Windows "open" is weird—like the pot calling the kettle white. Windows-based boxes are also no less fragmented than Android phones. Right here, I've got a Lenovo ThinkPad with proprietary UI features, such as the red point-stick and the ThinkVantage button. And, not every iPhone handset works the same. Old models either don't run or work well with the latest OS. Here's Steve again:

Twitter client, Twitter Deck, recently launched their app for Android. They reported that they had to contend with more than 100 different versions of Android software on 244 different handsets. The multiple hardware and software iterations present developers with a daunting challenge. Many Android apps work only on selected Android handsets running selected Android versions. And

this is for handsets that have been shipped less than 12 months ago.

Compare this with iPhone, where there are two versions of the software, the current and the most recent predecessor to test against.

Here's what TweetDeck (Steve got the name wrong) actually said in its blogged report (blog.tweetdeck.com/android-ecosystem):

To date, we've had 36,427 active beta testers, and below you can see the massive variety of phones and Android OS versions everyone is running. We were really shocked to see the number of custom roms, crazy phones and the general level of customization/hackalicious nature of Android. From our perspective, it's pretty cool to have our app work on such a wide variety of devices and Android OS variations.

So TweetDeck has no problem with "fragmentation". In fact, TweetDeck thinks it's cool. Also, what Steve's saying with his "two versions" point is that only the last two iPhone models are supported. Earlier ones are obsolete. In fact, they have to be, if you're keeping up with the latest downloads of the OS and its mother ship, iTunes. Now to the matter of smoke-blowing:

In reality, we think the open versus closed argument is just a smokescreen to try to hide the real issue, which is, what's best for the customer, fragmented versus integrated. We think Android is very, very fragmented and becoming more fragmented by the day. And as you know, Apple's provides the integrated model so that the user isn't forced to be the systems integrator.

...And we also think our developers can be more innovative if they can target a singular platform rather than a hundred variants. They can put their time into innovative new features rather than testing on hundreds of different handsets.

After hearing this, I contacted a friend who does lots of Android (and Linux) development. His first response:

For a phone OS, it's as open as they come, by a huge margin. Source code is readily available for free download by anybody. Having a complete copy of the source code makes development easier. But in general, you can do anything you want on Android.

Then he adds:

Here's where Steve's right: anybody can build their own Android device based on the resources available on-line. If you use Qualcomm chips, Qualcomm will have already written the kernel drivers for you and gotten it all integrated. Other chip vendors probably do the same thing. But if you want your device to work with the DRM-protected apps on the Android market, you have to have your device locked down security-wise and vetted by Google. Sometimes vendors push an over-the-air upgrade and forget to vet the new OS with Google first. Then all of a sudden paid apps stop working. Still, with Android you can generate a single binary that runs everywhere. Yes, there are version issues. Develop for 2.1, then 1.5 phones can't use it. That's why I always develop my apps on the most primitive version of the OS that will run them. The development environment (Eclipse) lets you do that. You can run any version of the OS under the emulator. And every app can specify a version of the OS it was intended for, plus a range of OS versions that it supports.

In the argument about arguments, it's easy to miss the Big Thing that's happening here: phones are becoming the new PCs, and there are many choices of ways to develop for them—and Android (and, therefore, Linux) is one of the top two. Apple deserves credit for cracking open the market with a highly vertical approach. But Apple's market is a closed silo. Android's market is wide open. There's no limit on how big it can be. ■

Doc Searls is Senior Editor of *Linux Journal*. He is also a fellow with the Berkman Center for Internet and Society at Harvard University and the Center for Information Technology and Society at UC Santa Barbara.

iX-Triton TwinBlade Servers: The Easy-to-Manage, Greener Way to Serve

► AFFORDABLE ► ECONOMICAL ► SAVINGS

The new **Triton TwinBlade Server** is the most technologically advanced blade server system in the industry, and the ideal solution for power-efficiency, density, and ease of management.



The **Triton TwinBlade Server** supports up to 120 DP servers with 240 Intel® Xeon® 5600/5500 series processors per 42U rack, achieving an unmatched 0.35U per DP node. Up to two 4x QDR (40 Gbps) Infiniband switches, 10GbE switches or pass-through modules give the TwinBlade the bandwidth to support the most demanding applications.

With **N+1 redundant, high efficiency (94%) 2500W power supplies**, the TwinBlade is the Greenest, most energy-efficient blade server in the industry. The energy saved by the iX-Triton TwinBlade Server will keep the environment cleaner and greener, while leaving the green in your bank account.

Server management is also simple with the Triton Twin Blade Server.

Remote access is available through SOL (Serial Over Lan), KVM, and KVM over IP technologies. A separate controller processor allows all of the Triton's remote management and monitoring to function regardless of system failures, offering true Lights Out Management.

Using the **Triton's management system, administrators can remotely control TwinBlades**, power supplies, cooling fans, and networking switches. Users may control the power remotely to reboot and reset the Triton TwinBlade Center and individual Twin Blades, and may also monitor temperatures, power status, fan speeds, and voltage.

For more information on the **iX-Triton TwinBlade**, or to request a quote, visit: <http://www.iXsystems.com/tritontwinblade>

Key features:

- Up to 10 dual-node TwinBlades in a 7U Chassis, 6 Chassis per 42U rack
- Remotely manage and monitor TwinBlades, power supplies, cooling fans, and networking switches
- Virtual Media Over Lan (Virtual USB, Floppy/CD, and Drive Redirection)
- Integrated IPMI 2.0 w/ remote KVM over LAN/IP
- Supports one hot-plug management module providing remote KVM and IPMI 2.0 functionalities
- Up to four N+1 redundant, hot-swap 2500W power supplies
- Up to 16 cooling fans

Each of the TwinBlade's two nodes features:

- Intel® Xeon® processor 5600/5500 series, with QPI up to 6.4 GT/s
- Intel® 5500 Chipset
- Up to 128GB DDR3 1333/ 1066/ 800MHz ECC Registered DIMM / 32GB Unbuffered DIMM
- Intel® 82576 Dual-Port Gigabit Ethernet
- 2 x 2.5" Hot-Plug SATA Drive Trays
- Integrated Matrox G200eW Graphics
- Mellanox ConnectX QDR InfiniBand 40Gbps or 10GbE support (Optional)



Call iXsystems toll free or visit our website today!
+1-800-820-BSDi | www.iXsystems.com



Cut Execution Time by >50% with WhisperStation-GPU

Delivered ready to run new GPU-enabled applications:

Design

3ds Max
Bunkspeed
Shot
Adobe CS5

Simulation

ANSYS Mechanical
Autodesk Moldflow
Mathematica

MATLAB
ACUSIM AcuSolve
Tech-X GPULib

BioTech

AMBER
GROMACS
NAMD, VMD
TeraChem

Integrating the latest CPUs with NVIDIA Tesla Fermi GPUs, Microway's WhisperStation-GPU delivers 2x-100x the performance of standard workstations. Providing explosive performance, yet quiet, it's custom designed for the power hungry applications you use. Take advantage of existing GPU applications or enable high performance with CUDA C/C++, PGI CUDA FORTRAN, or OpenCL compute kernels.

- ▶ Up to Four Tesla Fermi GPUs, each with: 448 cores, 6 GB GDDR5, 1 TFLOP single and 515 GFLOP double precision performance
- ▶ Up to 24 cores with the newest Intel and AMD Processors, 128 GB memory, 80 PLUS® certified power supply, and eight hard drives
- ▶ Nvidia Quadro for state of the art professional graphics and visualization
- ▶ Ultra-quiet fans, strategically placed baffles, and internal sound-proofing
- ▶ New: Microway CL-IDE™ for OpenCL programming on CPUs and GPUs



WhisperStation with 4 Tesla Fermi GPUs

Microway's Latest Servers for Dense Clustering

- ▶ 4P, 1U nodes with 48 CPU cores, 512 GB and QDR InfiniBand
- ▶ 2P, 1U nodes with 24 CPU cores, 2 Tesla GPUs and QDR InfiniBand
- ▶ 2U Twin² with 4 Hot-Swap MBs, each with 2 Processors + 256 GB
- ▶ 1U S2050 servers with 4 Tesla Fermi GPUs

Microway Puts YOU on the Cutting Edge

Design your next custom configuration with techs who speak HPC. Rely on our integration expertise for complete and thorough testing of your workstations, turnkey clusters and servers. Whether you need Linux or Windows, CUDA or OpenCL, we've been resolving the complicated issues – so you don't have to – since 1982.

Configure your next WhisperStation or Cluster today!

microway.com/quickquote or call 508-746-7341

Sign up for technical newsletters and special GPU promotions at microway.com/newsletter



OctoPuter™ 4U Server with up to 8 GPUs and 144 GB memory

1U Node with 2 Tesla Fermi GPUs

2U Twin² Node with 4 Hot-Swap Motherboards
Each with 2 CPUs and 256 GB



GSA Schedule
Contract Number:
GS-35F-0431N

Microway
Technology you can count on™