15 YEARS OF LINUX JOURNAL

KDE 4.2 | GNOME Do | Adobe AIR | Xfce | jQuery

LINUX JOURNAL

Since 1994: The Original Magazine of the Linux Community
MARCH 2009 | ISSUE 179

BROWSERS COMPARED
Opera, CrossOver Chromium and Flock

GNOME-Do
Launch Apps Quickly

THE DAY THE EARTH STOOD STILL
LINUX DOES SCI-FI

Xfce
An Alternative Lightweight Desktop

Create Your Own Virtual Reality with OpenSim

Check Out KDE's Latest Enhancements

PLUS:
Cross-Platform Apps with **Adobe AIR**

www.linuxjournal.com
$5.99US $5.99CAN
03
0 09281 03102 4

UBIQUITI NETWORKS

The Embedded Wireless Dream Machine.

680MHz CPU, 64MB RAM, 16MB Flash, USB, 3 mini-PCI slots, 5A power supply for multiple hi-power radios, and enhanced temperature performance and ESD protection for carrier applications. Full Linux SDK support and a $59 USD MSRP.

RouterStation

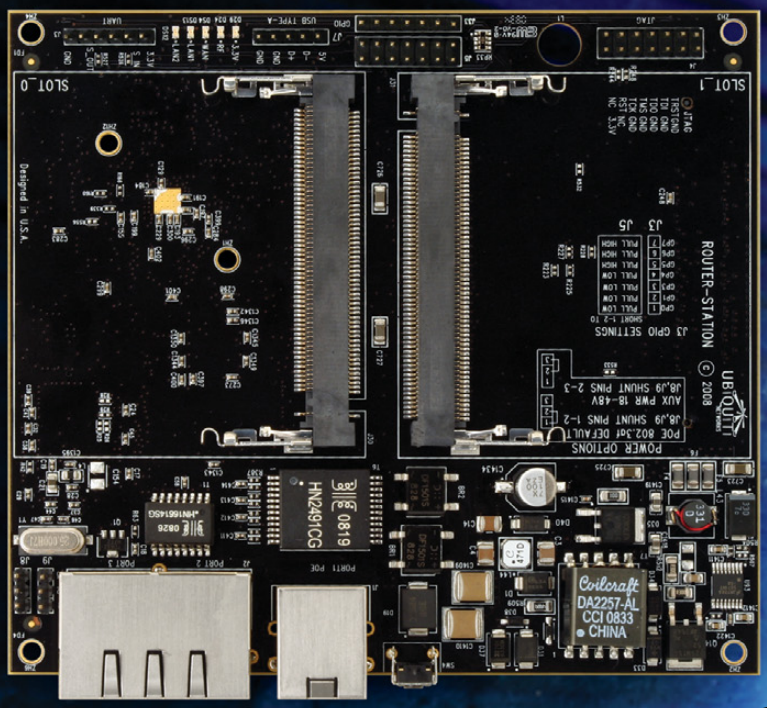

RouterStation UI CHALLENGE

CASH PRIZE

$200000

www.ubnt.com/challenge

www.ubnt.com

RouterStation UI Challenge: For Contest Details Visit: www.ubnt.com/challenge
Cash Prize $200000 (Two hundred Thousand) USD.

User Interface development based on OpenWrt Linux firmware. OpenWrt is a registered trademark of OpenWrt.

zt Systems

ZT Systems delivers something different: a unique

B A L A N C E

of world-class server performance and cost advantage joined with extensive flexibility in delivery and support

From the small office to the mega-datacenter, ZT Systems provides dependable, energy efficient server and client solutions that maximize productivity and optimize your IT budget. An established manufacturer celebrating our 15th anniversary, ZT combines individualized attention to your requirements with the experience and scalability you demand. All ZT Systems are assembled and serviced in the USA.

Choose ZT Systems and get:

- Custom Configurations
- Professional Grade Quality
- Personalized Account Service
- Superior Value

ZT Reliant SC1256i-40
Small Business Server
Entry-level Single Socket Solution

Starting at Only $749

ZT-SC1256i-40

- Intel® Core™2 Quad Processor Q8200
- 4GB DDR2 800MHz SDRAM (2 x 2GB)
- (2) 1.5TB SATAII 32MB Hotswap Hard Drives (Max. 4)
- 16X DVD-ROM Drive
- RedHat Enterprise 5 Trial Version
- 3 Year Limited Warranty and 24x7 Telephone support

ZT Systems 1105Ri-82
1U Rack Server
Entry-level file/web/email Server Solution

Starting at Only $949

ZT-1105Ri-82-C000001

- Intel® Xeon® Processor X3220
- 4GB DDR2 800MHz ECC Unbuffered SDRAM (2 x 2GB)
- (2) 500GB SATA Enterprise Hard Drives (Max. 4)
- 16X DVD-ROM Drive
- RedHat Enterprise 5 Trial Version
- 3 Year Limited Warranty and 24x7 Telephone support

ZT Systems 1213Ri-84
1U Rack Server
Data Center / Network Infrastructure and Cluster Solution

Starting at Only $1249

ZT-1213Ri-84-C000001

- Single Intel® Xeon® Processor E5410 (dual capable)
- 4GB DDR2 ECC/REG. SDRAM (2 x 2GB)
- (2) 500GB SATAII 32MB Hotswap Hard Drives (Max. 4)
- Slim DVD-ROM (Optional)
- RedHat Enterprise 5 Trial Version
- 3 Year Limited Warranty and 24x7 Telephone support

ZT Systems 4202Ci-84
4U/Tower Rack Server
Convertible Tower / Rack SMB Solution

Starting at Only $1649

ZT-4202Ci-84-C000001

- Single Intel® Xeon® Processor X5410 (dual capable)
- 8GB DDR2 ECC/REG. SDRAM (4 x 2GB, Dual Rank)
- (3) 1TB SATAII 32MB Hotswap Hard Drives (Max. 8)
- 16x DVD-ROM & 1.44MB Floppy Drive
- RedHat Enterprise 5 Trial Version
- 3 Year Limited Warranty and 24x7 Telephone support

Scalable Custom Server Solutions - Contact Us to Learn More

(866) 984-7687 corpsales@ztsystems.com

Intel, the Intel logo, Xeon, Intel Core, Xeon Inside and Core Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

CONTENTS

MARCH 2009
Issue 179

FEATURES

48 FAST APP LAUNCHING WITH GNOME DO

Use GNOME Do and its many plugins to extend desktop functionality.

Kyle Rankin

52 XFCE: THE THIRD MAN

The "other" desktop environment.

Federico Kereki

58 BROWSER BATTLES: OPERA, CROSSOVER CHROMIUM AND FLOCK

Stuck on Firefox? You may find fulfillment in an alternative browser.

James Gray

64 AIR OUT YOUR DESKTOP WITH ADOBE

Desktop apps for everyone.

Shawn Powers

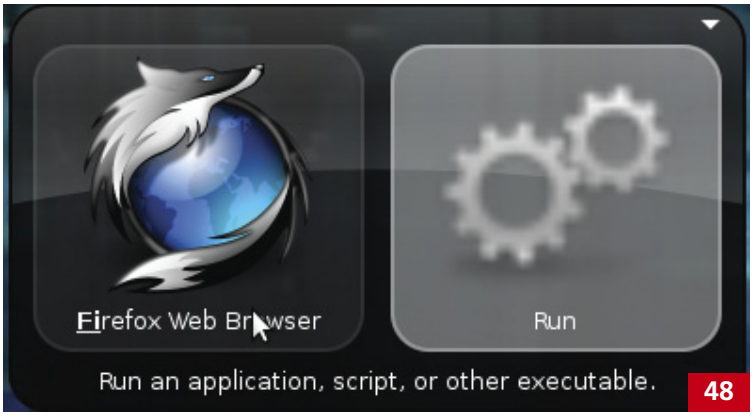

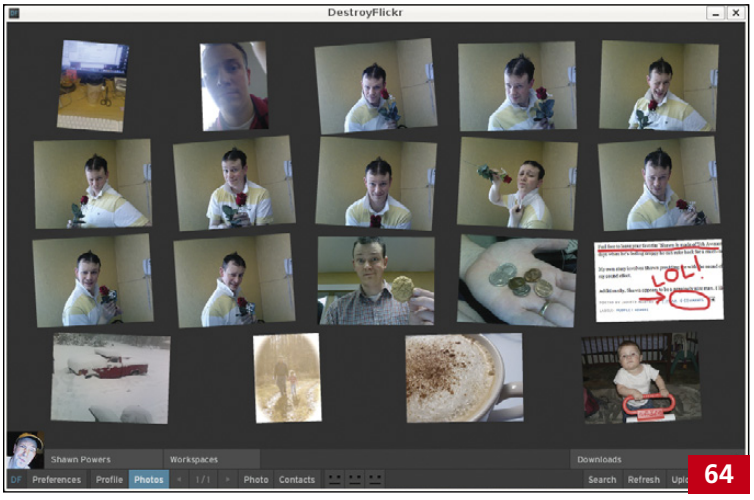

ON THE COVER

- **Browsers Compared: Opera, CrossOver Chromium and Flock, p. 58**
- **GNOME Do: Launch Apps Quickly, p. 48**
- ***The Day the Earth Stood Still*—Linux Does Sci-Fi, p. 72**
- **Xfce: An Alternative Lightweight Desktop, p. 52**
- **Check Out KDE's Latest Enhancements, p. 22**
- **Cross-Platform Apps with Adobe AIR, p. 64**

Cover image: TM and © 2008 Twentieth Century Fox Film Corporation.

servers DIRECT®

MORE PRODUCTS, BETTER SERVICE, GUARANTEED.

GO STRAIGHT TO THE SOURCE!

1.877.727.7887 | www.ServersDirect.com

YOUR HIGH PERFORMANCE COMPUTING HAS ARRIVED.

The ServersDirect® Systems with the Intel® Xeon® Processor helps you simplify computing operations, accelerate performance and accomplish more in less time

SOHO SERVER / STORAGE

OPTIMIZED FOR SMALL OFFICE & HOME OFFICE APPLICATIONS

SDP-S4100-T00

COST EFFECTIVE ENTRY LEVEL SERVER

- Mid Tower Chassis with 300W Power Supply
- Supermicro Server Board with Intel® 945GC Chipset
- 2x Realtek RTL8111C-GR Gigabit LAN Ports
- Support up to 4x SATA (3.0Gbps) Ports RAID 0, 1, 5, 10
- 4x 3.5" internal drive bays

STARTING AT **$329**

SDP-S4200-T00

MOST COST EFFECTIVE LAN SOLUTION

- Mid Tower Chassis with 300W Power Supply
- Supermicro Server Board with Intel® 3200 + ICH9R Chipset
- Support up to a Dual-Core Intel® Xeon® 3000 Series processor
- Support up to 6xSATA (3.0Gbps) Ports RAID 0, 1, 5, 10
- 4x 3.5" internal drive bays

STARTING AT **$439**

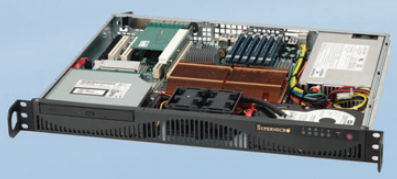

SDR-S1208-T00

STARTING AT **$559**

- Supermicro Mini 1U Rackmount Server with 260W Power Supply
- Supermicro Server Board w/Intel® 946GZ Chipset
- Support up to a Dual-Core Intel® Xeon® 3000 Series processor
- TPM Support
- 1x 3.5" Internal Drive Bay
- 2x Intel® 82573 PCI-e Gigabit LAN Port

SDR-C1303-T02

STARTING AT **$899**

- 1U Rackmount Server with 400W Power Supply
- Supermicro Server Board w/ Intel® 5100 Chipset
- Support Dual Intel® 64-bit Xeon® Quad-Core or Dual-Core
- Support up to 48GB 667/533MHZ DDR2 ECC Reg
- 2x 3.5" Hot-swap SATA Drive Bay
- Intel® 82573V & 82573L Gigabit Ethernet Controller

SDR-C2301-T06

STARTING AT **$999**

- 2U Rackmount Server with 460W Power Supply
- Supermicro Server Board w/ Intel® 5100 Chipset
- Support Dual Intel® 64-bit Xeon® Quad-Core or Dual-Core
- Support up to 48GB 667/533MHZ DDR2 ECC Reg
- 6x 3.5" Hot-swap SATA Drive Bay
- Intel® 82573V & 82573L Gigabit Ethernet Controller

SDR-C3301-T16

STARTING AT **$1,359**

- 3U Rackmount Server with 650W Power Supply
- Supermicro Server Board w/ Intel® 5100 Chipset
- Dual Intel® 64-bit Xeon® Quad-Core or Dual-Core
- Support up to 48GB 667/533MHZ DDR2 ECC Reg
- 16 x 3.5" Hot-swap SATA Drives Trays
- Dual-port Gigabit Ethernet Controller

SDR-C4302-T02

STARTING AT **$899**

- 4U Rackmount Server with 600W Power Supply
- Supermicro Server Board w/ Intel® 5100 Chipset
- Support Dual Intel® 64-bit Xeon® Quad-Core or Dual-Core
- Support up to 48GB 667/533MHZ DDR2 ECC Reg
- 3 x 3.5"Internal SATA Drives Trays
- Dual-port Gigabit Ethernet Controller

SDR-A8301-T42

STARTING AT **$4,799**

- 8U Chassis with 1350W Redundant Power Supply
- Supermicro Server Board w/ Intel® 5400 Chipset
- Support Dual Intel® 64-bit Xeon® Quad-Core or Dual-Core
- Support up to 64GB SDRAM Fully Buffered DIMM (FB-DIMM)
- 42 x 3.5"Internal SATA Drives Trays
- Dual-port Gigabit Ethernet Controller

SERVERS DIRECT CAN HELP YOU CONFIGURE YOUR NEXT HIGH PERFORMANCE SERVER SYSTEM - CALL US TODAY!

Our flexible on-line products configurator allows you to source a custom solution, or call and our product experts are standing by to help you assemble systems that require a little extra. Servers Direct - your direct source for scalable, cost effective server solutions.

1.877.727.7887 | www.ServersDirect.com

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, Pentium, and Pentium III Xeon are trademarks of Intel Corporation or it's subsidiaries in the United States and other countries.

CONTENTS MARCH 2009 Issue 179

COLUMNS

10 SHAWN POWERS' EDITORIAL
Celebrating 15 Years of *Linux Journal*

18 REUVEN M. LERNER'S AT THE FORGE
jQuery Plugins

22 MARCEL GAGNÉ'S COOKING WITH LINUX
The Evolution of the Desktop—How Far from the Pinnacle?

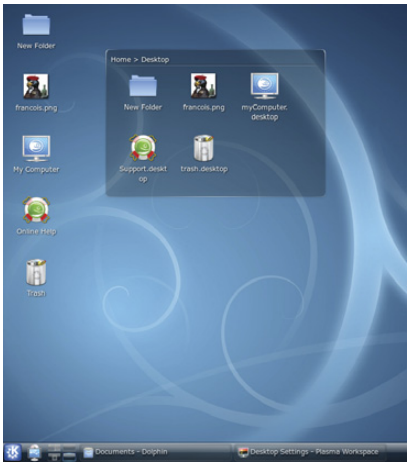

30 DAVE TAYLOR'S WORK THE SHELL
Counting Words and Letters

32 MICK BAUER'S PARANOID PENGUIN
Mental Laziness and Bad Dogma to Avoid

38 KYLE RANKIN'S HACK AND /
When Disaster Strikes: Hard Drive Crashes

80 DOC SEARLS' EOF
The Most Personal Device

INDEPTH

72 *THE DAY THE EARTH STOOD STILL*
Visual effects with Linux.
Robin Rowe

76 RUN YOUR OWN VIRTUAL REALITY WITH OPENSIM
Make your own reality, virtually.
Bill Childers

IN EVERY ISSUE

10 FROM THE EDITOR
12 LETTERS
14 UPFRONT
42 NEW PRODUCTS
44 NEW PROJECTS
81 ADVERTISERS INDEX

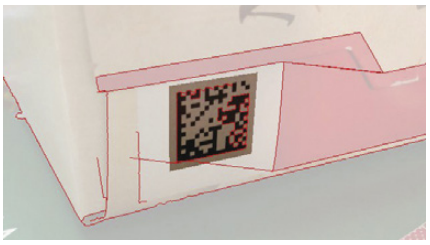

44 libdmtx

Next Month

SYSTEM ADMINISTRATION

They say a picture's worth a thousand words, so if you use Munin, you probably can stop talking for a few years. With Munin, you can graph anything you can monitor. And, if you like monitoring, we've got it covered, with monitoring from A to Zabbix.

Of course, there's not a lot to monitor if you can't get your system to boot, so check out FreeBoo. FreeBoo is a combination of open-source tools that provides network dual (multi) boot with all the power of fat clients and all the management ease of thin clients. If you use Rembo, FreeBoo can save you some serious coin.

With all the free time these tools will give you, you'll have time to deal with the really nasty problems, like bad Master Boot Records. But don't worry, Kyle Rankin's unofficial series "When Disaster Strikes" will help you through the next time you hose your MBR.

USPS *LINUX JOURNAL* (ISSN 1075-3583) (USPS 12854) is published monthly by Belltown Media, Inc., 2211 Norfolk, Ste 514, Houston, TX 77098 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 Bleuchip International, P.O. Box 25542, London, ON N6C 6B2

Polywell Linux Solutions

More Choices, Excellent Service, Great Value!

Serving the Industry for More Than 20 Years

Netdisk 8000V
Quiet Performance NAS Storage

The new generation of energy-efficient 45W AMD Athlon™ X2 dual-core processors with Cool'n'Quiet™ technology are designed to enable quiet performance.

4TB $1,399
8TB $2,399
12TB $2,999

- Dual Gigabit LAN
- RAID-5, 0, 1, 10
- Hot Swap, Hot Spare
- Linux, Windows, Mac
- E-mail Notification
- Tower or Rackmount

AMD Athlon™ X2 64

Silent Eco Green PC

Based on the new 45W AMD® Athlon™ X2 dual-core processor is energy efficient, quiet and has plenty of power. starts at **$199**

Fanless Silent ITX PC

1G DDR2, 80GB Hard Drive starts at **$299**
AMD® processor, Low-profile Add-on Available

AMD Opteron™ 64

4U 24Bay Storage Server

AMD Opteron™ Processors
36TB RAID-6 128GB RAM, 4 x GigaLAN

Small 1U Server for Data Center ISP

AMD Phenom™ processor
4 to 8GB DDR2-800, 2 x 500GB RAID HD
Linux Server Starts at **$499**

AMD Phenom™ 64

Polywell OEM Services, Your Virtual Manufacturer
Prototype Development with Linux/FreeBSD Support
Small Scale to Mass Production Manufacturing
Fulfillment, Shipping and RMA Repairs

- 20 Years of Customer Satisfaction
- 5-Year Warranty, Industry's Longest
- First Class Customer Service

888.765.9686
linuxsales@polywell.com
www.polywell.com/us/Lx

Polywell Computers, Inc 1461 San Mateo Ave. South San Francisco, CA 94080 650.583.7222 Fax: 650.583.1974

AMD, AMD Athlon, Phenom, Opteron, Cool'n'Quiet, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other names are for informational purposes only and may be trademarks of their respective owners.

POLYWELL

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/digital

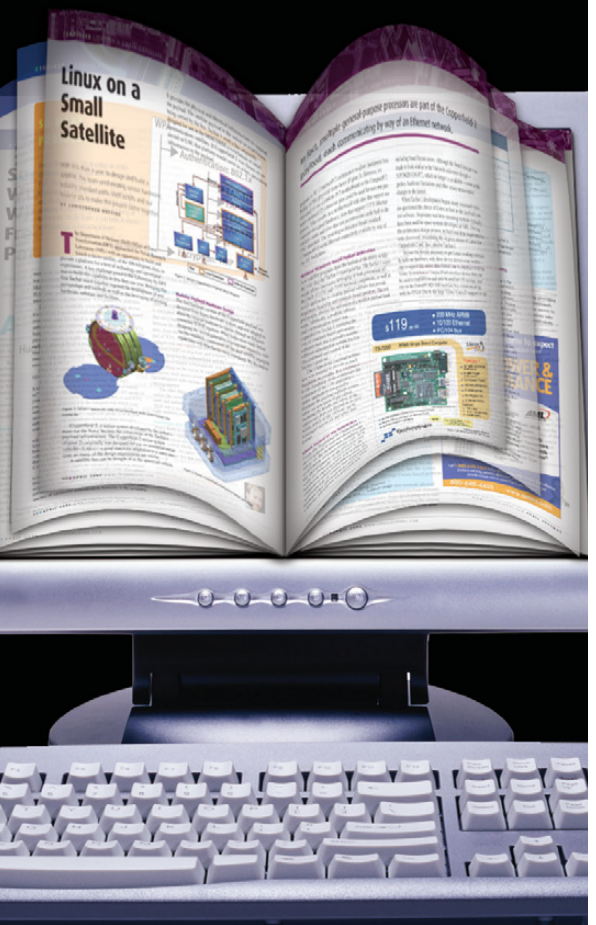

LINUX JOURNAL

Executive Editor	Jill Franklin jill@linuxjournal.com
Senior Editor	Doc Searls doc@linuxjournal.com
Associate Editor	Shawn Powers shawn@linuxjournal.com
Associate Editor	Mitch Frazier mitch@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
Chef Français	Marcel Gagné mggagne@salmar.com
Security Editor	Mick Bauer mick@visi.com
Hack Editor	Kyle Rankin lj@greenfly.net

Contributing Editors
David A. Bandel • Ibrahim Haddad • Robert Love • Zack Brown • Dave Phillips • Marco Fioretti
Ludovic Marcotte • Paul Barry • Paul McKenney • Dave Taylor • Dirk Elmendorf

Proofreader Geri Gale

Publisher	Carlie Fairchild publisher@linuxjournal.com
General Manager	Rebecca Cassity rebecca@linuxjournal.com
Sales Manager	Joseph Krack joseph@linuxjournal.com
Sales and Marketing Coordinator	Tracy Manford tracy@linuxjournal.com
Circulation Director	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

Reader Advisory Panel
Brad Abram Baillio • Nick Baronian • Hari Boukis • Caleb S. Cullen • Steve Case
Kalyana Krishna Chadalavada • Keir Davis • Adam M. Dutko • Michael Eager • Nick Faltys • Ken Firestone
Dennis Franklin Frey • Victor Gregorio • Kristian Erik • Hermansen • Philip Jacob • Jay Kruizenga
David A. Lane • Steve Marquez • Dave McAllister • Craig Oda • Rob Orsini • Jeffrey D. Parent
Wayne D. Powel • Shawn Powers • Mike Roberts • Draciron Smith • Chris D. Stark • Patrick Swartz

Editorial Advisory Board
Daniel Frye, Director, IBM Linux Technology Center
Jon "maddog" Hall, President, Linux International
Lawrence Lessig, Professor of Law, Stanford University
Ransom Love, Director of Strategic Relationships, Family and Church History Department,
Church of Jesus Christ of Latter-day Saints
Sam Ockman
Bruce Perens
Bdale Garbee, Linux CTO, HP
Danese Cooper, Open Source Diva, Intel Corporation

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.

PRINTED WITH SOY INK®

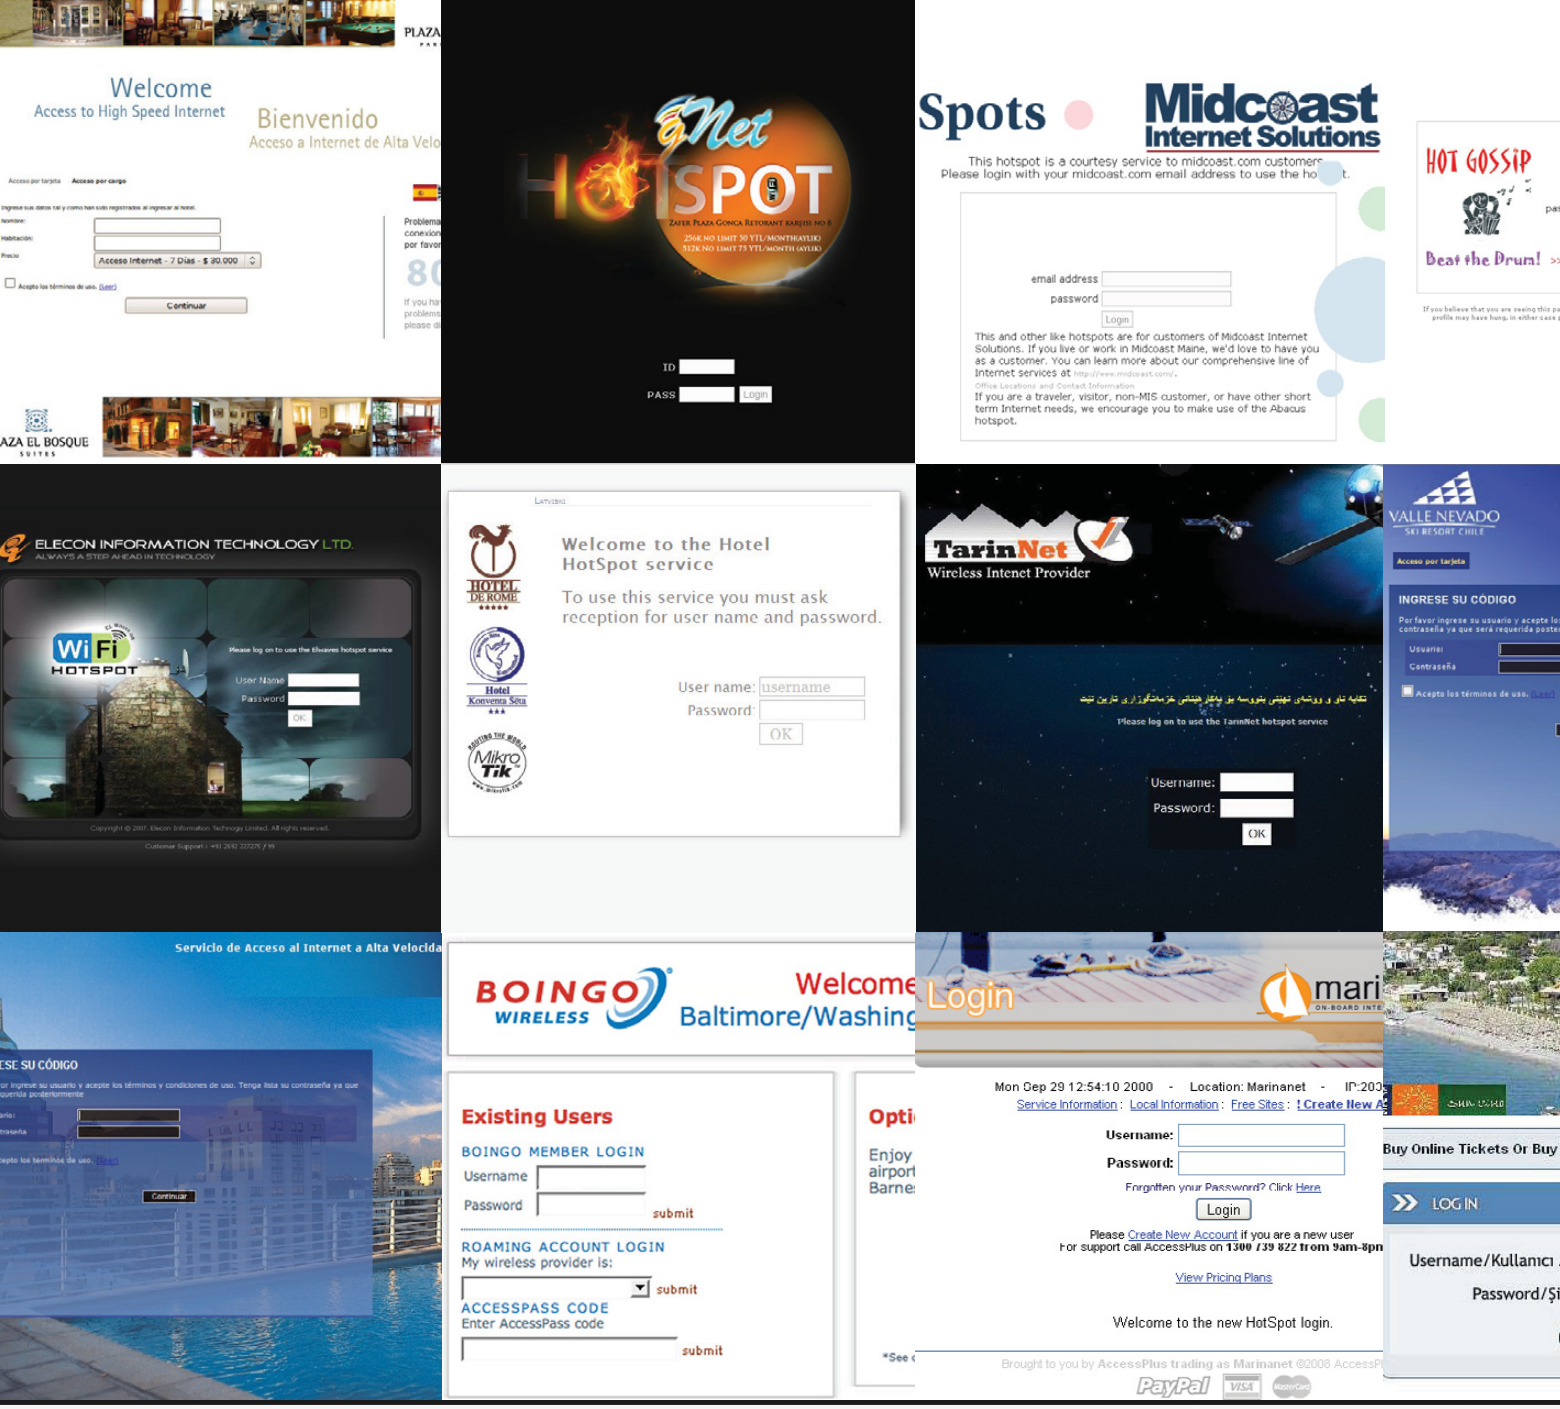

MikroTik powered Hotspots around the world

MikroTik RouterOS powered hotspots are everywhere. From Internet cafes in desert towns of Africa, to futuristic airport lounges in the US and five star hotels in the Mediterranean. Mikrotik can power your systems too. Free evaluation installations are available in our download section.

www.mikrotik.com

CELEBRATING 15 YEARS OF *LINUX JOURNAL*

See the first issue of *LJ*, published in March 1994, at **www.linuxjournal.com/issue/1**.

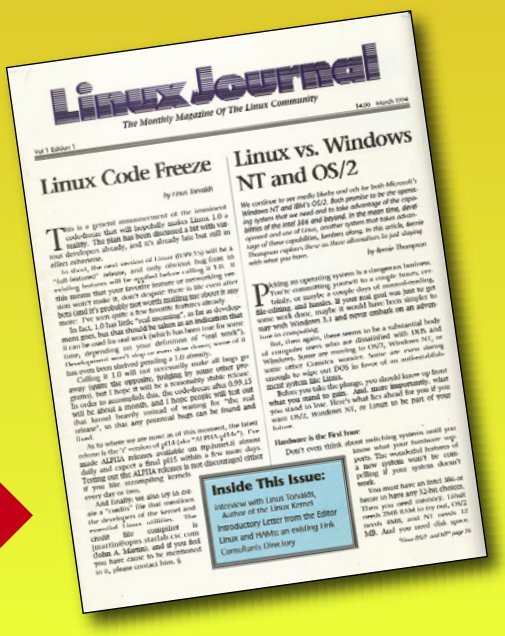

SHAWN POWERS

[Shawn Powers was barely out of high school when the first issue of Linux Journal went to press 15 years ago, but we figured it would be fun to send him back in time to write a column for the first issue. Besides, how else could we claim a tax write-off on a time machine?]

Wow, what a dream. I could have sworn I was a 30-something-year-old geek with a family and a mortgage. But, here I am in 1994. Oh well, at least the *Linux Journal* gig wasn't a dream. What's *Linux Journal*, you ask? That's easy. We're the only magazine dedicated to the Linux and Open Source community. What's "Open Source"? Well, you'll have to wait a few years for that one.

Our publisher, Bob Young, brings us a great interview with Linus Torvalds. As I'm sure you know, Linus has quite a bit to do with the Linux community. Bob Young also is someone you'll want to keep an eye on—trust me on this one, and maybe think about investing in red-colored headwear. We've also got an awesome comparison of the three leading operating systems: Linux, Windows and OS X—er, I mean OS/2. Sure, IBM is pouring a ton of money into marketing its operating system, while Linux doesn't have any huge financial backing, but I think history will prove it takes more than hype to compete. Looking through the vista of time, Microsoft will have its fair share of blunders too. Linux is here to stay.

The big news that comes along with this maiden issue of *Linux Journal* is that the Linux kernel itself has matured to 1.0 status. Just because it's no longer beta doesn't mean you'll have to start paying for it though. Linux is free—free in several ways. Check out Arnold Robbin's "What's GNU?" column for more details.

Are you worried you won't be able to run Linux on your existing hardware? Well, admittedly, hardware compatibility is a challenge, but if you're looking to install a basic Linux system, you should expect to have a computer with at least 2MB of RAM and 15MB of disk space. Also, the fancy 386 processor will give you amazing 32-bit speeds and fully utilize the power of the Linux kernel. With the 386 math coprocessor and its 32-bit address space, I can't imagine we'll ever need a faster processor. ::SNORT::

One of the biggest announcements this month is the availability of a new Linux distribution called Debian. Ian Murdock, the creator and maintainer of Debian, tells us why his distribution is different and why the Linux community needs something like it. He has the backing of the Free Software Foundation and is making the entire operating system available as a free download to anyone who wants it—awesome stuff that will almost certainly stand the test of time. Again, trust me.

The one thing I'm sad to report is that in order to try all the awesome programs available for Linux, you'll either have to download them very slowly from FTP servers or spend some money buying CDs. Installing from CDs is much faster though, so it might be worth the investment. I'd give you the Ubuntu CD I brought with me, but I fear it might disrupt the space-time continuum. And, I probably would get in trouble for that.

I'm so excited for everyone stuck back here in 1994. You have years and years of *Linux Journal* issues to read. Whatever your current plans are for the Linux operating system, keep subscribing to *Linux Journal*, and we'll keep you up to date with the latest information, tech tips, programming practices and industry news for the next 15 years and beyond!

Now, if you'll excuse me, I need to go find Linus. He hasn't figured out what sort of mascot to choose for Linux, and I'm a big penguin fan. If I start trying to convince him now, maybe in a few years, he'll decide my penguin idea is a good one. Wish me luck!■

In 1994, Shawn was attending his first year of college at Michigan Tech University. He skipped his engineering classes almost every day to sneak into the computer labs and play with Linux. At the time it seemed a waste of tuition, but looking back, he wouldn't change a thing.

EmperorLinux

...where Linux & laptops converge

Powerful Linux: The Rhino

- Based on the Dell Precision M6400/Latitude E6500
- High performance NVidia 3-D on a WUXGA widescreen
- High performance Core 2 Quad, 16 GB RAM
- Ultimate configurability — choose your laptop's features

Features include:

- 2.2-3.0 GHz Core 2 Duo/Extreme (dual-core) or 2.5 GHz Core 2 Quad QX9300 (quad-core)
- Up to 17" WUXGA LCD w/ X@1920x1200
- NVidia Quadro FX 3700M graphics, 128 core CUDA
- 80-320 GB hard drive (7200 rpm SATA) or 128 GB solid state drive, up to 16 GB RAM
- DVD±RW or Blu-ray, Ethernet, 802.11a/g/n, Bluetooth
- One year Linux tech support - phone and email
- Three year manufacturer's on-site warranty
- Choice of pre-installed Linux distribution:

Portable Linux

Toucan T500/W500

- ThinkPad T500/W500 by Lenovo
- Up to 15.4" WUXGA w/ X@1920x1200
- ATI Radeon graphics
- 2.2-2.8 GHz Core 2 Duo
- Up to 4 GB RAM
- 100-320 GB hard drive / 128 GB SSD
- 5.3-6.5 pounds
- Starts at $1280

Tablet Linux

Raven X200 Tablet

- ThinkPad X200 tablet by Lenovo
- 12.1" WXGA w/ X@1280x800
- 1.2-1.86 GHz Core 2 Duo
- Up to 4 GB RAM
- 80-320 GB hard drive / 128 GB SSD
- Pen/stylus input to screen
- Dynamic screen rotation
- Starts at $2065

Rugged Linux

Tarantula CF-30

- Panasonic Toughbook CF-30
- Fully rugged MIL-SPEC-810F tested: drops, dust, moisture, & more
- 13.3" XGA TouchScreen
- 1.6 GHz Core 2 Duo
- Up to 4 GB RAM
- 80-320 GB hard drive
- Call for quote

www.EmperorLinux.com 1-888-651-6686

Model prices, specifications, and availability may vary. All trademarks are the property of their respective owners.

letters

Clarification from Damballa

This letter is in regard to Grzegorz Landecki's "Detecting Botnets" article from the January 2009 issue of Linux Journal*, from Paul Royal, Principal Researcher at Damballa, Inc.—Ed.*

As the individual who discovered the Kraken botnet, I wanted to clarify some inaccuracies that were stated regarding its discovery and handling. In addition to these clarifications, I would welcome any conversations about Damballa, its processes and Kraken findings if Mr Landecki is interested in further discussion.

From the *Linux Journal* article:

"To simplify, Damballa discovered (probably during a security audit) a new malware with hard-coded addresses (URLs) for zombie control servers (CCs—computers that manage tasks for zombie machines and all infected computers report to them)."

Damballa first discovered Kraken through network anomalies observed on a DynDNS provider's network with which it collaborates to shut down C&C servers. Damballa originally thought Kraken was HacDef (a well-known type of malware used by a botnet documented in 2006; see **www.avira.com/en/threats/section/fulldetails/id_vir/0/bds_hacdef.fv.1.a.html**) and did not have a malware sample for Kraken until late March 2008. Damballa only discovered that the botnet was not HacDef after it located a malware sample matching the network anomaly. That others in industry thought the botnet was Bobax, an older botnet (and perhaps ancestor of Kraken), speaks to the general difficulty of new botnet identification.

From the article: "Damballa also found that some of those hard-coded addresses were not registered in a DNS service (the botnet probably was tested at that time, and the authors were preparing to launch it later)."

Kraken malware does not use hard-coded domain names at all; it autogenerates them. The algorithm used to generate domain names has been documented by both ThreatExpert (**blog.threatexpert.com/2008/04/kraken-changes-tactics.html**) and TippingPoint (**dvlabs.tippingpoint.com/blog/2008/04/28/owning-kraken-zombies**).

From the article: "Damballa registered those domains as its own and ended up controlling quite a large botnet for research."

Damballa never "controlled" Kraken; because Kraken malware generates domains, the botmasters simply registered domains that would be subsequently autogenerated (Damballa could not register them all) and resumed its operation.

From the article: "It hasn't contacted any security company about the methods of infection it discovered."

In early April 2008, Damballa released

Tech Tip Tuning

I like to read the Tech Tips, as I did with David S. Sinck's tip "Treating Compressed and Uncompressed Data Sources the Same" (*LJ*, December 2008, page 56). Here are some remarks and suggestions.

I strongly recommend using double quotes. The presented solution silently misses all files with one or more whitespaces in its name.

The only purpose of the line:

```
F=$(echo $1 | perl -pe 's/.gz$//')
```

is to remove the last file extension from the filename. This is done with Perl, an interpreter for a language much more powerful than Bash, a subshell and a pipe. There's too much effort to remove a few characters (in every loop cycle!). The first two lines of the function could be replaced by a parameter substitution, which costs nearly nothing. So, with quoting, we get the following:

```
function data_source ()
{
  local F=${1%.gz}            # remove file extension

  if [[ -f $F ]] ; then
    cat "${F}"
  elif [[ -f $F.gz ]] ; then
    gunzip --stdout "${F}.gz"
  fi
}    # -------  end of function data_source  -------

for file in * ; do
  data_source "$file" | ...
done
```

--
Fritz Mehner

instructions for remediating Kraken compromises (**www.damballa.com/downloads/brochures/Kraken_Remediation.pdf**) and a technical whitepaper/industry response (**www.damballa.com/downloads/press/Kraken_Response.pdf**) that included Kraken C&C domains that companies/organizations could use to prevent compromised systems from talking to the Kraken control servers.

Damballa (who codeveloped and administrates the malfease malware repository—**www.malfease.oarci.net**) also shared malware samples (that could be downloaded by researchers/practitioners) with SANS ISC (**isc.sans.org/diary.html?storyid=4256**) and OffensiveComputing (**www.offensivecomputing.net/?q=node/699**).

From the article: "It hasn't published any details of the exploits used to any bugtrack, nor has it contacted any vendors to alert them of the issue."

Kraken didn't use any exploit to propagate; there were no vendors to contact. Like Storm, Kraken propagated through social engineering (e-mail and instant message).

--
Paul Royal

PHOTO OF THE MONTH

Have a photo you'd like to share with *LJ* readers? Send your submission to publisher@linuxjournal.com. If we run yours in the magazine, we'll send you a free T-shirt.

Tux and Tuxina cake, submitted by David Backeberg, taken at his friend's wedding.

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**.

7" Touch Panel Computer
for embedded GUI / HMI applications

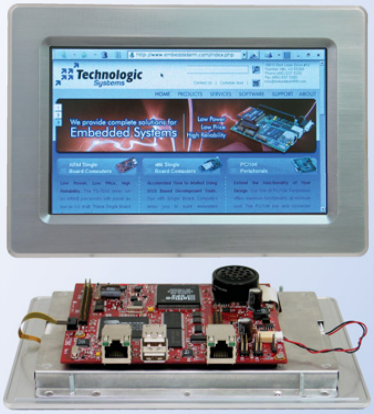

quantity 1 pricing starts at $449

Powered by a
200 MHz ARM9 CPU

- Low power, Industrial Quality Design
- Mountable aluminum frame
- 64MB SDRAM (128MB opt)
- 512MB Flash w/ Debian Linux
- Programmable FPGA - 5K LUT
- 7" Color TFT-LCD Touch-Screen
- 800x480 customizable video core
- Dedicated framebuffer - 8MB RAM
- Audio codec with speaker
- Boots Linux 2.6 in about 1 second
- Unbrickable, boots from SD or NAND
- Runs X Windows GUI applications
- Runs Eclipse IDE out-of-the-box

Our engineers can customize for your LCD

- Over 20 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

See our website for our complete product line

Technologic SYSTEMS

We use our stuff.
visit our TS-7800 powered website at
www.embeddedARM.com
(480) 837-5200

diff -u

WHAT'S NEW IN KERNEL DEVELOPMENT

Much of the Linux kernel source tree is taken up with the text of **licenses** under which the code is released. There are plenty of GPL-compatible licenses, and much of the kernel is released under a dual license. As **Jörn Engel** points out, a lot of kernel source files also include the full text of their licenses, causing about 3.5MB of duplication throughout the source tree, or something more than 1% of the total size of the tree. **Roland Kletzing** recently reported those numbers, suggesting that one way to reduce the size of the kernel would be to replace all that duplication with simple one-line pointers to a single copy of each relevant license.

Although completely rational, this may not be likely to happen. On the one hand, as Jörn says, a lot of the duplication is done because corporate legal departments want their engineers including full licensing information with all their contributions. Taking it out again might pose problems for them. On the other hand, **Chris Snook** reminded folks that the kernel was distributed in compressed form, and because the licenses were all duplications, they compressed down to a very, very small size. **Bodo Eggert**'s objection—that even this small size, when multiplied by the number of times the kernel actually was downloaded, resulted in an immense waste of bandwidth—was largely ignored, primarily because shrinking the compressed sources by such a small amount also could be accomplished by doing actual coding work, which is more fun, more sexy, and has the added benefit of actually making the kernel run better.

Maybe someone will find a way to keep the corporate lawyers happy and reduce the amount of license duplication at the same time, but this doesn't seem likely to become a high priority for anybody.

There's a new driver out that connects PCI devices called **PCINet**. It was written by **Ira Snyder**, and it implements the Ethernet protocol across the PCI interface. Now folks can code up communication between two PCI devices, using the familiar Ethernet protocol.

No one likes getting a **kernel oops**. Well, some folks actually do, and they spend lots of time trying to invoke them, just so they can fix them and make the world better. **Arjan van de Ven** has written a Perl script to take the raw oops output, pretty it up, and also try to clarify the part of the oops that's actually relevant to figuring out what went wrong. The script, called **markup_oops.pl**, still is tethered to the i386 platform, thanks to some coding constructs that **Keith Owens** pointed out to him. It also seems that Perl is no longer the "language of ascendancy", so writing the script in Perl may be another strike against it.

Michael Halcrow, **Tyler Hicks** and **David Kleikamp** have added filename encryption to **eCryptFS**. This has some interesting issues that make the implementation trickier than it might seem. For example, the encryption process can make the filename slightly longer than the unencrypted version, so if the unencrypted version is too close to the maximum length allowed by the filesystem, the encryption process won't work.

Also, the eCryptFS folks want encryption to be optional—some files in the directory may be encrypted while others aren't. How do you tell the difference between an encrypted filename and a filename that's just really dorky? The answer is that the encrypted names all have a special prefix, and the rest of the name is in a special format. If you really want to mess with the filesystem and create an unencrypted filename with that prefix and that format, you're on your own. But the issue probably will never come up.

The **linux-next tree** is continuing to be fine-tuned. On the one hand, **Stephen Rothwell** wants the tree to be a place where code migrates quickly into the main kernel tree. He also specifically wants **Andrew Morton** to be able to base his own -mm tree releases on linux-next easily. All of this means linux-next is intended to make sure that code merges well together and will not cause housekeeping problems for any upstream maintainers (such as **Linus Torvalds** and Andrew).

So, all code is expected to be unit-tested and posted to the relevant mailing lists. The folks submitting the code should know that if it doesn't get dropped, it'll be heading into the official tree. Any patches that conflict with the tree will be dropped. Any code that breaks the build will be dropped. In practice, truly trivial breaks won't get code dropped from linux-next, if Stephen can just eyeball the problem and fix it quickly. But, anything beyond that will require the folks submitting the code to fix it up and submit it again. In some cases, people from different projects will have to work together to resolve whatever breakage one of their submissions brought to light. The idea is that linux-next gives these projects the opportunity to identify and work on those conflicts and other problems, before trying to get it past Linus or Andrew and into their trees.

—ZACK BROWN

LJ Index, March 2009

1. Average number of minutes it takes to read an on-line privacy policy: 10
2. Mean word count of privacy policies for the Web's most popular sites: 2,514
3. Number of words in MySpace's Terms and Conditions: 3,868
4. Number of words in the MySpace Terms and Conditions paragraph that says, "It is therefore important that you review this Agreement regularly to ensure you are updated as to any changes.": 65
5. Number of Web sites visited by the average person per month: 119
6. Number of hours it would take to read all privacy policies for 119 Web sites: 20
7. Millions of Web users in the US: 121.667
8. Billions of hours per year that would be spent if every US Web user read all Web site privacy policies: 44.3
9. Estimated time cost in dollars per year, per US Web user, for reading all Web site privacy policies: 3,000
10. Estimated billions of dollars in lost leisure and productivity time that would be cost by all US Web site users reading all privacy policies: 365
11. Millions of Web sites in Google's first index, in 1998: 26
12. Billions of pages indexed by Google in 2000: 1
13. Trillions of pages indexed by Google in 2009: 1
14. Year in which lines of code in the Linux kernel will pass 6.5 million: 2009
15. Millions of dollars it would have cost "by conventional proprietary means in the US" to have produced Red Hat Linux 6.2 in 2000: 600
16. Billions of dollars it would have cost to produce Red Hat Linux 7.1 in 2002: 1
17. Those same billions in 2008 dollars: 1.2
18. Billions of dollars it would have cost to produce Fedora 9 in 2008: 10.8
19. Billions of dollars it would have cost to produce the Linux kernel alone in 2008: 1.4
20. Estimated billions of dollars spent by Microsoft on Vista client operating software, Windows Server 2007 server software, Office 2007, Internet Explorer 7 and other related software: 14

Sources: 1, 2, 5–10: *Ars Technica*, Carnegie Mellon CyLab, Out-Law.com, all sourcing "Cost of Reading Privacy Policies" by Carnegie Mellon researchers Aleecia McDonald and Lorrie Faith Cranor | 3, 4: MySpace 11–13: Jesse Alpert and Nissan Hajaj, Software Engineers, Google Web Search Infrastructure Team 14: Greg Kroah-Hartman | 15, 16: David A. Wheeler 17–19: Amanda McPherson, Brian Proffitt and Ron Hale-Evans of the Linux Foundation | 20: Dennis Byron

Fast Is Hot

Last month, we reported here about Splashtop (**www.splashtop.com**), which starts a laptop in only a few seconds. Since then, Splashtop reportedly has found its way into the Lenovo IdeaPad S10e and the ASUS Eee PC.

Now there's Cloud (**thinkgos.com/press-release20081201.php**), from Good OS, which gave us gOS Linux and a cover story for *Linux Journal* in March 2008. Good OS calls Cloud "A New Operating System for 2009". More specifically, the press release says Cloud "integrates a Web browser with a compressed Linux operating system kernel for immediate access to Internet, integration of browser and rich-client applications, and full control of the computer from inside the browser" (**thinkgos.com/press-release20081201.php**).

The browser looks like Google's Chrome, and most of the icons in the Mac-like dock on the bottom of the screen are for Google apps running in "the cloud". One exception is the Windows symbol. We won't go there. Meanwhile, the fast-start race all goes to Linux.Cloud, launched in December 2008 at the Netbook World Summit in Paris, demonstrated on a GIGABYTE Netbook.—DOC SEARLS

Find More at LinuxJournal.com

This month's issue is all about Linux on the Desktop, and we have plenty more where this came from at LinuxJournal.com.

How better to learn some new desktop tricks than checking out some of our videos at **www.linuxjournal.com/video**?

Elliot Isaacson demonstrates how to run your Windows partition without rebooting: **www.linuxjournal.com/video/run-your-windows-partition-without-rebooting**.

You also may enjoy reading about Adobe's Linux offerings for the Adobe AIR Runtime: "An Introduction to AIR": **www.linuxjournal.com/content/introduction-air**.

Have you ever wanted to run an old DOS favorite on your Linux desktop? Check out Mitch Frazier's article "Running DOS Programs on Linux: Duke Nukem Lives!": **www.linuxjournal.com/content/running-dos-programs-linux-duke-nukem-lives**.

As always, Bruce Byfield brings a constant flow of new tips and tricks for using OpenOffice.org: **www.linuxjournal.com/content/openofficeorg-many-views-impress**.

Visit us at LinuxJournal.com often! There's always a new tip, trick or how-to to keep you busy.—KATHERINE DRUCKMAN

They Said It

An avidity to punish is always dangerous to liberty. It leads men to stretch, to misinterpret, and to misapply even the best of laws. He that would make his own liberty secure must guard even his enemy from oppression; for if he violates his duty he establishes a precedent that will reach to himself.

—Thomas Paine,
www.brainyquote.com/quotes/authors/t/thomas_paine.html

In the longer perspective, Linux will become a serious alternative for our high-end phones.

—Nokia VP Ukko Lappalainen,
www.reuters.com/article/internetNews/idUSTRE4B16IO20081202

Coincident with the release of this code I have concluded my agreement with Atheros whereby I had access to information about their devices. This means that in the future all fixes, updates for new chips, etc., will need to be a community effort. Atheros states the Linux platform will be the reference public code base, so folks wanting to add support for other platforms will have to scrape the information from there.

—Sam Leffler, maintainer of the binary HAL, also used in MadWifi releases, on release of source code for his HAL (hardware abstraction layer) variant, lwn.net/Articles/308984

Expect to see millions of Web devices, even desktop Web devices, in the coming years that completely strip out the Windows layer and use the browser as the only operating system the user needs.

—Mike Arrington,
www.techcrunch.com/2008/09/01/meet-chrome-googles-windows-killer/comment-page-2

WiMAXing Linux

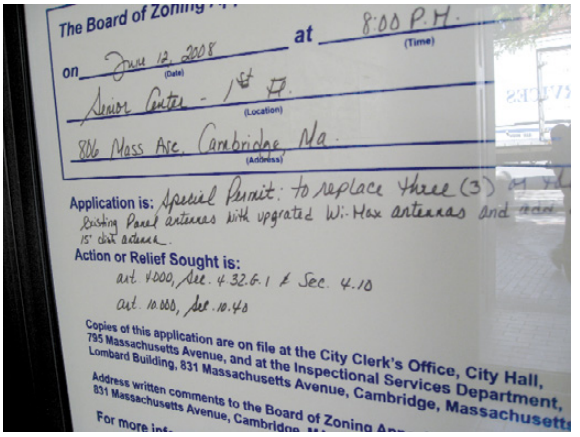

Connect these dots:

- After going through a complicated merger with Sprint-Nextel's wireless broadband operation, Clearwire expects to add headcount (in a down economy!) as it rolls out WiMAX services across the country (**seattletimes.nwsource.com/html/businesstechnology/2008457502_brier02.html**). The new entity will be called Clear.

- Clear expects "Mobile WiMAX" download speeds to reach "10 to 15 Mbps" within two years.

- The merger (**gigaom.com/2008/11/28/the-iphone-and-the-ensuing-wireless-broadband-boom**) was supported by an investment of $3.2 billion by Bright House Networks, Comcast, Google, Intel and Time Warner Cable.

- Google wishes to see Linux-based Android devices work on the network. Says *Ars Technica* (**arstechnica.com/news.ars/post/20081201-future-of-wimax-is-clear-as-sprint-clearwire-close-deal.html**), "Even though Google partners with just a subset of the industry, the availability of unencumbered devices with Google services requires competitive moves toward openness from other carriers." The keyword is "unencumbered".

- Intel, maker of WiMAX chipsets, has been investing heavily in WiMAX device development and network deployment. Recently, it joined the Ministry of Economic Affairs (MOEA) in Taiwan to jointly establish an "enabling center" for Moblin (**moblin.org**), which calls itself "a Linux-based platform optimized for the next generation of mobile devices, including Netbooks, Mobile Internet Devices and In-vehicle infotainment systems". On that last front, Intel is working with Wind River on a Moblin-enabled "automotive infotainment stack". Says LinuxDevices, "Wind River hopes the latter will appear in products from 'Tier 1' automotive equipment makers, and it lists Bosch, BMW Group, Delphi and Magneti Marelli among the companies 'actively supporting' its efforts to 'drive Linux', so to speak, into the automotive infotainment market." Intel Capital is also putting $11.5 million into Taiwan's VMAX Telecom, which also plans to roll out WiMAX data services.

- Linpus (**www.linpus.com**), also based in Taiwan, is joining the Moblin movement. Linpus is behind the Linux-based Acer Aspire One Netbook. Chetan Sharma Consulting (**www.slideshare.net/chetansharma/us-wireless-market-q3-2008-update-nov-2008-chetan-sharma-consulting-presentation**) says the ARPU (average revenue per user) for voice in the US is on a downward trend, and "data ARPU will become a more dominant factor of the overall APRU mix by the end of 2009. The percentage contribution from data is likely to exceed 25% by the end of 2008 and 30% by the end of 2009." Sourcing a report by Analysis Mason, Gigaom reports (**gigaom.com/2008/11/28/the-iphone-and-the-ensuing-wireless-broadband-boom**) that "By 2015, data will account for 94% of total wireless network traffic in developed regions."

- *BusinessWeek* reports (**www.businessweek.com/magazine/content/08_49/b4111064905299.htm**) that Netbooks are "cannibalizing PC and laptop sales".

That's just a small sampling of the rumblings where Linux and mobile data meet. But the photo above also shows how WiMAX cells are actually being rolled out. For more on the future of mobile data developments, see "The Most Personal Device", on page 80.

—DOC SEARLS

Looking for a domain?

More ways to get the domain name you REALLY want!

New!

Register new domains …

When you search for a domain name, our domain checker automatically gives you relevant domain suggestions, ensuring that you find the best domain name for your website.

… or buy a domain that has already been registered.

If the domain name you want has already been registered, you can use our FREE domain auction service to search over 14 million domain names and make a bid on the one you want.

FREE Private Domain Registration

Your privacy is important to us. With free private domain registration, your name, address and phone number remain confidential. Stop spammers, telemarketers, data miners and anyone else who might misuse your personal data.

Save 25% on .com domains!*

Offer ends February 28th!

*Offer valid through February 28, 2009. Discount applied to first year of a 2 year minimum contract term. Visit www.1and1.com for full promotional offer details. Product and program specifications, availability and prices 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.
© 2009 1&1 Internet, Inc. All rights reserved.

MEMBER OF united internet

Call **1-877-GO-1AND1**

Visit us now **www.1and1.com**

1&1

What They're Using

Don Stokes

When I started following Linux action on Twitter, one contributor stood out: mich_linux_guy. His tweets include, "Playing with the Android SDK", "Getting ready to install Fedora 10 on my son's XO laptop....Hope it's a win" and "Bought an Acer Aspire One at Costco for $350. Just installed Ubuntu 8.10. Did require fiddling with wireless packages (Googled it). Works!" His name is Don Stokes. He's a veteran hardware and software hacker living with his wife and three sons in Oakland County, Michigan, where he's also pursuing an MS in Computer Science at Oakland University. His startup is **stokesautomation.com**. And his favorite quote is from Spinoza: "Let every man think what he wants and say what he thinks." What follows is plenty of both.

Don's Four-Year-Old Son on His OLPC

The first computer I purchased was a TRS-80 model 1. I goofed around with some text-based adventure games and played with BASIC programming. The next computer I bought had an Intel 80286 processor. It came with MSDOS, but I bought a copy of IBM OS/2 to run. Eventually, OS/2 died, and I was forced to run Windows NT. I think it was in 1999 that a friend at work (Ben—what happened to you? Call me!) passed me a Red Hat Linux CD. (It was actually labeled "Red Ass Linux".) It required a lot of tinkering to get it to work, but for me, this was "fun"! Eventually, I decided to make it my primary OS.

For our anniversary, I bought my wife an HP laptop with a gig of RAM and put Ubuntu on it. She uses it for her e-mail, Web browsing and to manage the music on her iPod. I put Ubuntu on my dad's old Dell laptop about a year ago. As he doesn't have the root password or sudo rights, he can't mess it up. Because it is a lot harder to attack Linux than Windows, I don't worry about the security of his system. It's a great OS for him. My mom is next!

I ran Fedora until Ubuntu Feisty Fawn came out. Once in a while, I will boot a live CD of other distributions to see what they are up to. I also have Damn Small Linux (DSL) running on a couple of machines. It is amazing how much functionality you can get using this distribution on machines that most people would leave on the curb for the garbage truck.

The best thing about Linux is that its modular architecture makes it so configurable/customizable. If I want eye candy, I can run GNOME or KDE desktops. I can even run the X server with only a display manager, if I want to maximize efficiency. On machines with low CPU or memory resources, I can run with only a command line on a text display with no graphical user interface.

Another reason I use Linux is the huge repository of free software. Kids have very short attention spans. Even the best games won't keep them content for long. With Linux, I can simply go to the Add/Remove option on the main menu and choose from a gazillion free applications—all installable with a few mouse clicks. Many of these applications are very educational; it's been a great experience for my kids. Because they don't have administrator access, they can't corrupt the machine.

I recently downloaded the Android SDK for the Eclipse IDE. It is amazing how easy it is to create applications written in Java for Android. I got so excited, I went out and bought a G1 the next day. (That's how long it took the local T-Mobile guys to get some stock.) I'm currently writing a simple encryption application for storing private information like passwords and account numbers. Collecting devices that run Linux has become kind of an obsession with my family (okay, me). Here is our current collection:

- HP Pavilion zv5000 laptop with Ubuntu 8.10 (my primary development PC).
- HP Pavilion dv6000 laptop with Ubuntu 8.10 (wife's kitchen PC).
- HP Pavilion ze4200 laptop with Xubuntu 8.10 (games and KTurtle for my eight-year-old son).
- HP Pavilion a1687 desktop with Ubuntu 8.10 (hooked to our LCD TV in the family room).
- ASUS Eee PC running Xubuntu 8.10 (rides in my bag).
- Acer Aspire One running Ubuntu 8.10 (maybe a replacement for the Eee?).
- OLPC XO running Fedora, Sugar and XFCE (toy for my four-year-old son—he loves hearing it say what he types).
- IBM NetVista with 128MB of RAM running Xubuntu 8.10.
- Compaq Pressario, 166MHz Pentium MMX, 16MB of RAM running Damn Small Linux (basement machine—with the graphical desktop up playing a CD, it is using only 8MB!).
- Six other PCs of varying CPU and RAM running Ubuntu and DSL (acting as file/print/game servers).
- Nokia n810 (replacement for my Palm TX, but replaced by G1).
- T-Mobile G1 phone running Android.

All the software I use on a daily basis is free, and almost all of it is open source. I use OpenOffice.org for word processing, spreadsheets, presentations and simple databases. I use Mozilla Firefox and Thunderbird for Web browsing and e-mail. Pidgin is my instant messenger. I use gcc for compiling my C/C++ source and Eclipse for my Java development. I also use Apache Web servers and a MySQL database manager for systems I have developed. I also use Umbrello, Octave, Kino and VLC on a regular basis. I love my freedom!

—DOC SEARLS

Need a server for your small business?

Look no further.

Introducing 1&1 Premium Servers, a powerful new line of dedicated servers specifically designed for high performance needs. Featuring energy efficient technology, these top-of-the-line machines reduce costs and environmental impact with increased performance-per-watt.

Special offers end February 28th!

All 1&1 Premium Servers include*:

- RAID
- Plesk 8 with 100-domain license
- FTP backup space
- Serial Console: Connect directly to the serial port of your server, even if the network is down.
- Recovery Tool: Load rescue image and reboot server.
- GeoTrust Dedicated SSL certificate (a $49 value!)
- Easy-to-configure firewall and more!

1&1 BUSINESS SERVER II

Dual-Core AMD Opteron™ 1218, 2 x 2.6 GHz,
4 GB DDR RAM, 2 x 500 GB Hard Drive

~~$199.99~~ per month

3 months FREE!*

HURRY! Special offer ends February 28th.

1&1 ENTERPRISE SERVER I

Quad-Core AMD Opteron™ 1352, 4 x 2.1 GHz,
4 GB DDR RAM, 2 x 750 GB Hard Drive

~~$299.99~~ per month

3 months FREE!*

HURRY! Special offer ends February 28th.

1&1 ENTERPRISE SERVER II

Quad-Core AMD Opteron™ 1356, 4 x 2.3 GHz,
8 GB DDR RAM, 2 x 1,000 GB Hard Drive

~~$399.99~~ per month

3 months FREE!*

HURRY! Special offer ends February 28th.

* Offer valid through February 28, 2009. Prices based on Linux servers. Discount applied to first 3 months of a 12 month minimum contract term, $99 setup fee applies. See www.1and1.com for full promotional offer details. Certain features not available with Managed Servers. Product and program specifications, availability and prices 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.

© 2009 1&1 Internet, Inc. All rights reserved.

MEMBER OF united internet

Call **1-877-GO-1AND1**

Visit us now **www.1and1.com**

1&1

REUVEN M. LERNER

jQuery Plugins

Learn how to use jQuery plugins, one of the secrets to this JavaScript library's success.

Last month, we began to look at jQuery, an open-source JavaScript library that provides a great deal of functionality for Web developers, which is increasingly popular for client-side application development. We saw that jQuery's use of CSS-style selectors, combined with its "chaining" syntax, makes it easy to get started with the library and to attach behaviors to page elements. We also saw that jQuery inherently is unobtrusive, with event handlers being assigned via `$(document).ready()`, rather than inline with the HTML.

At the end of the day though, jQuery does many of the same things as Prototype, YUI and other JavaScript libraries. So, why have so many developers moved to jQuery? What makes it such an attractive choice? Speed and the API are obviously two factors, but a major reason for developers to use jQuery is the huge library of plugins that is available for it. Just as Perl programmers can enjoy a massive library of modules known as CPAN, jQuery users can benefit from a large number of plugins for a variety of tasks, from UI elements to AJAX form submission. Installing and using a jQuery plugin is extremely straightforward, and it can be installed (and evaluated) in minutes.

This month, let's look at a few of the many jQuery plugins that have been developed over the last few years, and also at how to use plugins to change our Web applications.

Plugin Basics

From a developer's perspective, a jQuery plugin is nothing more than an additional JavaScript file that you download, install in your Web application's JavaScript directory, include in your program with a <script> tag and then invoke. Typically, a plugin adds one or more new functions to the jQuery object, which means if you install a plugin named foo, you often can do the following:

```
$(document).ready(function() {
    $('#mybutton').foo();
}
```

The above construct tells jQuery that when the HTML document has been downloaded enough to start querying and modifying it with JavaScript, you invoke a function. That function, in turn, looks for an HTML element with the ID mybutton and then invokes the foo() method on it.

What does `$('#mybutton').foo()` do? That's up to the author of the plugin. Typically, a plugin adds functionality to an element or class of elements, quite possibly modifying the HTML around that element—adding new elements necessary for the plugin to do its job or adding classes that cause one or more CSS declarations to be invoked.

Because a jQuery plugin typically modifies the document's HTML, it's vitally important to look at a plugin's documentation to understand what HTML structure it expects to receive. Perhaps it expects to have an unordered list (<ul>) with list items (<li>) inside it. Perhaps it expects to have tags with tags inside it. Perhaps it expects something else altogether. If a plugin doesn't seem to do what you expected, double-check that your HTML matches the example and/or what's in the documentation.

jQuery plugins also rely in no small part on the powerful visualizations that CSS provides. Installing a jQuery plugin often means not only using JavaScript code, but also putting CSS styles into effect—either by incorporating the plugin's CSS file into your application or by copying the declarations into an existing CSS file. Just as many plugins require that your HTML be structured a certain way in order to work, some require that you set certain classes or IDs on your HTML elements.

The fact that jQuery plugins modify the HTML and/or CSS means that you might need some extra tools to understand and debug what is happening in your browser. I normally develop in Firefox, and I have found the Firebug extension to be a wonderful tool to identify issues and experiment with alternatives, in both JavaScript and CSS. Also quite valuable is the Web Developer extension for Firefox, whose "view generated HTML" does the same thing as "view source", but shows you the HTML as it currently exists, not as it was downloaded originally from the server before JavaScript modified it.

Finally, some plugins come with images that enhance the way the plugin works.

The way I've described it so far might make it seem as though jQuery plugins are difficult to use or that they will force you to change and contort your HTML in numerous ways. But, nothing

is further from the truth. On the contrary, the main problem I have had with jQuery plugins is the vast variety and scope of them, forcing me to choose among 20 different menuing plugins or ten different modal-dialog plugins. Comparing and evaluating these plugins, many of which have borrowed code from one another, can be difficult. But, when you find an appropriate plugin, it's usually quite easy to get started using it.

If you are trying to do something in jQuery that others probably have tried before, you always should look through the main plugin repository first, at **plugins.jquery.com**. There also is a large library of user-interface (UI) plugins at **ui.jquery.com**. And of course, large numbers of plugins are described, documented and downloadable from Web sites outside jquery.com.

DataTables

HTML tables have been around for many years, and although they have gotten a bad reputation because of the way they were used and abused for layout purposes (even after the introduction of CSS), there are many times when a table is the best and most logical way to present data. If you are running an on-line store, for example, and you want to get a summary of recent orders, it makes sense to structure the data in a table.

One of the most common things users want to do with a table, once they see it, is sort the rows according to one particular column. To continue with our e-commerce example, perhaps they want to sort the list by order number. Or, perhaps they want to sort it by customer name, by price or by date.

It's not hard to do this kind of sorting on the server side. Set up the table headers to be clickable links, and when you get a request, you change the order of the rows before they are output. But, if the data already is in your browser, wouldn't it be nice to be able to sort the rows in JavaScript? This might not be the fastest possible way to execute such a sort, but given small enough data sets, it's acceptable for most purposes, and it gives the user a sense of desktop-like control and response.

One nice jQuery plugin I've found to do this is called DataTable, written by Allan Jardine (see Resources). DataTable takes an existing HTML table and makes it sortable by column, as well as searchable.

To get this to work, you need an HTML table. Listing 1 is an HTML file that will work just fine for these purposes, although you presumably will want to use DataTable with a dynamic Web application.

Listing 1. table.html

```
<html>
 <head>
  <title>Testing tables</title>
  <script type="text/javascript" src="jquery.js"></script>
  <script type="text/javascript" src="jquery.dataTables.js"></script>
  <script type="text/javascript">
    $(document).ready(function () {
        $('#people-table').dataTable();
      });
  </script>
 </head>

 <body>
  <h1>Testing tables</h1>
  <table id="people-table">
    <thead>
      <tr>
        <th>ID</th>
        <th>Last Name</th>
        <th>First Name</th>
        <th>City</th>
        <th>Balance</th>
      </tr>
    </thead>
    <tbody>
      <tr>
        <td>1</td>
        <td>Lerner</td>
        <td>Reuven</td>
        <td>Modi'in</td>
        <td>100</td>
      </tr>
      <tr>
        <td>2</td>
        <td>Barack</td>
        <td>Obama</td>
        <td>Washington</td>
        <td>750000000</td>
      </tr>
      <tr>
        <td>3</td>
        <td>Bush</td>
        <td>George</td>
        <td>Dallas</td>
        <td>-1000000000</td>
      </tr>
    </tbody>
  </table>
 </body>
</html>
```

As you can see, Listing 1 contains a single table with an ID of people-table. The table is defined as you might expect for an HTML table, with one possible exception (unless you're extremely pedantic). The headers for the table are defined with a <thead> section, while the body is in a <tbody> section. These tags are optional according to the HTML standard, but they are mandatory if you want to use DataTable.

With the table in place, you now can add jQuery and the DataTable plugin. Unlike other sorts of plugins, there's nothing to install, except the JavaScript file itself. If you put jquery.js and dataTable.js in the same directory as the file (which is probably not a good idea on a production system), you can write:

```
<script type="text/javascript" src="jquery.js"></script>
<script type="text/javascript" src="jquery.dataTables.js"></script>
```

Now everything is ready, except one thing. You need to connect the DataTable plugin to your table. You do this by telling JavaScript that when the document is ready, you want to connect the two:

```
$(document).ready(function () {
        $('#people-table').dataTable();
    });
```

If you aren't familiar with jQuery already, you'll soon learn that this is a common idiom when using the library. Define a function that is triggered on `$(document).ready` and have that function execute and/or define a number of other items, each of which fires on a different tag, ID or class.

Once you connect DataTable to the table, you'll see that the headers are now clickable and force the table to be sorted, first in ascending and then in descending order. (The DataTable download includes icons and CSS appropriate for seeing the sort order.)

DataTable supports a large number of options, all of which are passed to the dataTable() function. For example, DataTable shows ten rows of the current table by default, but it lets you choose from 10, 25, 50 and 100 rows. If you prefer to see a different number, you can set the iDisplayLength property to a different default:

```
$(document).ready(function () {
        $('#people-table').dataTable({'iDisplayLength': 1});
    });
```

I recently have used DataTable in a few projects, and I've found it to be easy to use, stable and well documented. The biggest problems crop up when you have a large data set, but that's not unique to DataTable.

Menus

Another common task people want to do in JavaScript is produce menus, including hierarchical menus. Indeed, if I think back several years, menus probably are one of the things for which my clients have most commonly asked. One of the best-known methods for creating menus with CSS is known as Suckerfish, because of the sample data that was used to show the technique.

There are many jQuery-based menu libraries, but one I've grown to enjoy is called Superfish, because it adds functionality to the Suckerfish style of menu.

Listing 2. menu.html

```
<html>
 <head>
  <title>Testing menus</title>
  <link rel="stylesheet" type="text/css" media="screen"
        href="superfish.css" />
  <script type="text/javascript" src="jquery.js"></script>
  <script type="text/javascript" src="superfish.js"></script>
  <script type="text/javascript">
    $(document).ready(function () {
        jQuery('ul.sf-menu').superfish();
      });
  </script>
 </head>

 <body>
  <h1>Testing menus</h1>

  <div>
    <ul class="sf-menu sf-navbar">
      <li>
        <a href="">Account</a>
        <ul>
          <li><a id="checking-menu" href="#checking-div">Checking</a>
          </li>
          <li><a id="savings-menu" href="#savings-div">Savings</a>
          </li>
          <li><a id="credit-card-menu" href="#credit-card-div">
                    Credit card
              </a>
          </li>
        </ul>
      </li>
      <li><a id="profile-menu" href="#profile-div">Profile</a>
      </li>
      <li><a id="help-menu" href="#help-div">Help</a>
      </li>
    </ul>
  </div>
 </body>
</html>
```

It handles submenus, adds shadows and even tries to be intelligent about when you plan to open the menu and when your mouse is passing by, using a separate plugin known as hoverIntent.

To use Superfish, you need to download and install the plugin. Then, you need to create a menu using a combination of <ul>, <li> and <a> tags. If you need a secondary hierarchy of menus, you can create one with a nested <ul> in an <li> tag. In each <a> tag, the href identifies which div on a page should be displayed when that menu item is clicked on, hiding all of the other divs by default.

You undoubtedly will want to start off with the Superfish CSS file that comes with the plugin. You always can modify it to suit your needs. There are large numbers of definitions, and I've never been able to build the CSS file from scratch. Instead, I've modified the existing one, changing it to suit my needs.

As always in jQuery, you use the plugin by attaching it to an element of the HTML page. Instead of using the element's ID, as you did with DataTable, here you attach it to the <ul> tag with a class of sf-menu:

```
<script type="text/javascript">
  $(document).ready(function () {
          jQuery('ul.sf-menu').superfish();
      });
</script>
```

If there were more than one <ul> with that class, Superfish would create menus on all of them. Remember, the jQuery object can return any number of page elements: zero, one or a large number. The full HTML for the example is shown in Listing 2.

Conclusion

Plugins are the secret to jQuery's success, and there are so many plugins for jQuery, it's impossible to describe them all here. But, as you can see from these two examples, using the plugin often requires very little effort. Once you get the hang of it, downloading, installing and using plugins becomes second nature. I've found it can be useful to create a simple, small HTML file with dummy data and use a jQuery plugin with that, just to understand the basics of how to use a plugin.

There are times when plugins clash with one another, in that they're both trying to rewrite the HTML, sometimes in conflicting ways. For example, I recently used DataTable along with with a jQuery tab widget, and it took me a while until I could ensure that everything was visible on the page. As jQuery plugins become increasingly sophisticated, we might have to worry about this more and more.

For now, however, jQuery plugins are a fun and easy way to spruce up your Web application. Next month, I'll explain how to design your own plugin, delving a bit deeper into jQuery's plumbing and understanding how jQuery takes advantage of JavaScript's quirks to give us an extensible platform for client-side programs.■

Resources

The jQuery home page is at **www.jquery.com**, and it includes a large number of links to tutorials and articles about the library.

The home page for the DataTables plugin is **www.sprymedia.co.uk/article/DataTables**.

The original article introducing CSS Suckerfish drop-down menus is at **www.alistapart.com/articles/dropdowns**. The Superfish jQuery plugin is at **www.alistapart.com/articles/dropdowns**.

If you are a Ruby on Rails developer and are interested in using jQuery (instead of the default Prototype and Script.aculo.us), you can learn more by reading **errtheblog.com/posts/73-the-jskinny-on-jquery**.

Reuven M. Lerner, a longtime Web/database developer and consultant, is a PhD candidate in learning sciences at Northwestern University, studying on-line learning communities. He recently returned (with his wife and three children) to their home in Modi'in, Israel, after four years in the Chicago area.

System on Module
Internet Appliance Engine
• Atmel ARM9 400Mhz CPU
• 6 Serial Ports & 2 SPIs
• Up to 60 Digital GPIOs
• 10/100 BaseT Ethernet
• SODIMM Bus Expansion
• SSC/I2S Audio Interface
• Linux with Eclipse IDE
• Real Time Clock Calendar
• SD/MMC Flash Card Interface
• Up to 512 MB Flash & 64 MB RAM
• 2 USB 2.0 Host Ports & 1 Device Port
• 4 10-Bit A/Ds & 6 16-Bit Timer/Counters
2.6 KERNEL
Microsoft Windows CE
The SoM-9G20 uses the same small SODIMM form-factor utilized by other EMAC SoM modules, and is the ideal processor engine for your next design. All of the ARM9 processor core is included on this tiny board including: Flash, Memory, Serial Ports, Ethernet, I2S Audio Interface, PWMs, Timer/Counters, A/D, digital I/O lines, Clock/Calendar, and more. Like other modules in EMAC's SoM product line, the SoM-9G20 is designed to plug into a custom or off-the-shelf Carrier board containing all the connectors and any additional I/O components that may be required. The SoM approach provides the flexibility of a fully customized product at a greatly reduced cost. Single unit pricing starts at $150.
http://www.emacinc.com/som/som9G20.htm
Since 1985 OVER 24 YEARS OF SINGLE BOARD SOLUTIONS
EMAC, inc.
EQUIPMENT MONITOR AND CONTROL
Phone: (618) 529-4525 • Fax: (618) 457-0110 • Web: www.emacinc.com

MARCEL GAGNÉ

The Evolution of the Desktop—How Far from the Pinnacle?

Rumor has it that 2009 is the year of the Linux desktop. With KDE 4.2, the next step in the evolution of the popular desktop environment, it may well be.

That is very impressive, François! I wasn't expecting to see a series of murals when I walked in, but there they are. What exactly am I looking at? You want me to guess, do you? Well, if I *had* to guess, I would say it looks like a series of tables from the prehistoric to the present, with place settings to match. Thank you, François. I particularly love that stone-age table with the bone utensils, but I still don't get the point. What made you do this? The evolution of the tabletop theme of today's menu? *Non, non, non*, François. That's evolution of the *desktop*, as in a Linux computer desktop.

François, don't look so sad. It's still quite an impressive oeuvre. Besides, our guests are arriving and there's no time to change things. Quickly, put a smile on your face and get ready.

Welcome, *mes amis*, to *Chez Marcel*, the home where fine wine is paired with delectable open-source software. Please sit and make yourselves comfortable. While you find your tables, perhaps François could make his way to the wine cellar. Check the north wing, *mon ami*, and you'll find a few cases of 2005 Gessinger Zeltinger Sonnenuhr Riesling from Germany. *Vite!*

There have been numerous improvements and enhancements in this new version of KDE, too many to cover them all in the space I have, so I'll give you a sampler of what you can expect. Think of it as a KDE 4.2 buffet table.

Everyone (well, almost everyone) loves fancy desktop effects. Flash and pizzazz are the order of the day with modern desktops, and KDE 4.2 doesn't lack there. To help with the seemingly mandatory collection of desktop special effects, KDE 4.2 now detects your graphic card's capabilities. If the card supports the compositing effects, they are turned on automatically.

Let's tour this new desktop and take a look at a few of the more interesting changes. In the December 2008 issue of *Linux Journal*, I told you about those cool little desktop gadgets, or widgets, called plasmoids. By way of a quick recap, KDE's new desktop isn't so much a desktop as a shell that runs other applications. It's called Plasma. Plasma runs small applications called plasmoids, though they also are referred to as widgets and even gadgets. Each one of those plasmoids is, like the Plasma shell, a container that can contain other plasmoids. Imagine turtles standing on top of one another, and you'll start to get the idea. One of KDE 4.2's enhancements is the sheer number of widgets it can run and the different types of widgets that it includes.

Click the plasma cashew icon in the top right-hand corner of your desktop, and select Add Widgets. When the Add Widgets window appears (Figure 1), you can select widgets included with KDE 4.2. To install from other sources (including Mac OS X dashboard widgets, among others), click the Install New Widgets button, and then choose Install Widget from Local File.

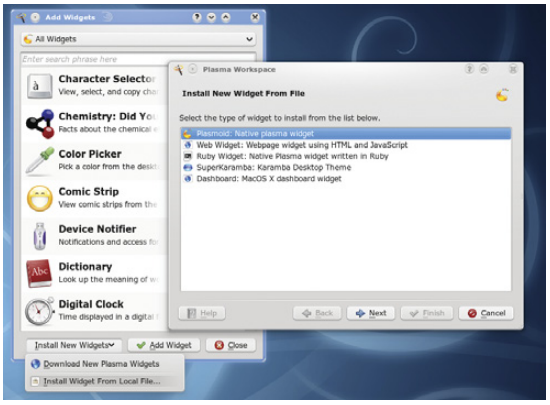

Figure 1. KDE 4.2 lets you install widgets from an amazing number of sources.

While we are on the subject of plasmoids, I must tell you about what was probably the most controversial change to accompany KDE 4.0: the so-called missing desktop icons. This became a controversy because desktop icons are simply files and folders in

Figure 2. Your old desktop icons, safely ensconced in the Folder View plasmoid.

a special folder called Desktop. To display the icons on the desktop (or in the Desktop folder), you would need a plasmoid whose entire job in life was to display the contents of the Desktop folder.

A storm ensued. I personally don't like icons on my desktop, but apparently, many (if not most) people do, and the KDE team heard about it—a lot. To calm this storm, they created the plasmoid I just mentioned—the one that would display the contents of the Desktop folder (or any folder for that matter). It's called Folder View (Figure 2).

Your plasma desktop can have multiple configurations to help you with whatever work you might happen to be doing. These are called activities, and you can add new ones by clicking that cashew-like icon in the top right-hand corner of your desktop (Figure 3). Two of these activities are defined for you by default. One provides the, ahem, classic desktop view with the optional Folder View plasmoid. The other is called Folder View. It provides the kind of desktop view with which most people are familiar, one where you can right-click to add icons and shortcuts to programs, files or URLs.

To switch from one defined activity to the other, left-click on the plasma cashew icon in the top right, and select Appearance Settings. When the dialog appears, you can select a desktop theme, change your wallpaper and, yes, switch activities from desktop to Folder View (Figure 4). Incidentally, running your desktop session in Folder View mode doesn't preclude running the Folder View plasmoid. Both can coexist quite happily, as you can see in Figure 5. That's because the Folder View plasmoid isn't only for your Desktop folder. It could be

Expert included.

Johnny is a team leader in our hardware production facility, and he's responsible for operating system installations and validation testing. He is one of the experts who makes certain that when you use the Silicon Mechanics online configuration tools to order your server, you receive exactly what you specified, correctly configured and ready to perform. Johnny sized up the new 1U server from Silicon Mechanics—the Rackform nServ A108—and he's excited about its energy efficiency and cost effectiveness.

The AMD Opteron™ 1000 Series processors available with this server are offered in low power consumption models, and the A108 boasts an 86% efficient power supply. With a starting configuration price below $750, the Rackform nServ A108 makes an outstanding entry-level server.

When you partner with Silicon Mechanics, you get more than an energy-efficient and cost-effective server configured just the way you want it—you get an expert like Johnny.

For more information about the Rackform nServ A108 visit www.siliconmechanics.com/A108.

SILICON MECHANICS

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

Silicon Mechanics and the Silicon Mechanics logo are registered trademarks of Silicon Mechanics, Inc.
AMD, the AMD Arrow logo, AMD Opteron, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

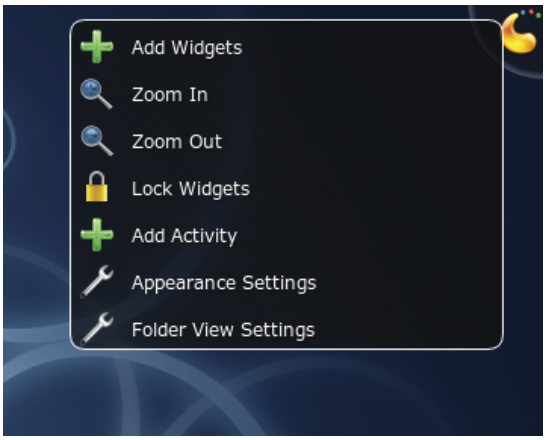

Figure 3. That little plasma cashew does so much, from changing your desktop theme to adding widgets (plasmoids), and even switching between—and adding—different activities.

Figure 4. Getting your cake and eating it—icons on the desktop and a Folder View plasmoid live together in perfect harmony.

anything you like.

All of this talk of activities is really just a means to keep things organized in a way that makes some kind of sense. In the bottom left-hand corner of the panel, you'll see a big K button. That's KDE's program launcher, called Kickoff (Figure 6). Kickoff may seem a bit alien at first, but you're going to love it. There are five animated tabs at the bottom of Kickoff window. Favorites is just what it sounds like—a list of favorite applications. By default, you'll see a small handful as provided by your distribution. Under the Applications tab, applications are listed under groups, such as Office, Games and so on.

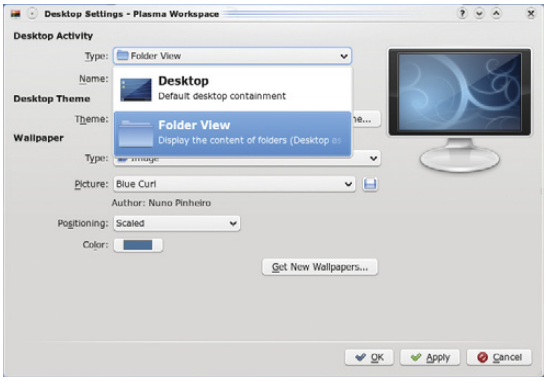

Figure 5. From the Desktop Settings dialog, you can switch desktop activities, along with the wallpaper, paint and so on.

Click on a program group, and the window slides to the next level of application until you get a list of the programs themselves. If you regularly traverse the office menus to find OpenOffice.org's Impress program, you might want to add it to your Favorites tab. Simply right-click on the application, and select Add to Favorites.

Conversely, you also can remove applications from the Favorites list in a similar way. Simply right-click, and select Remove from Favorites. Under the Computer tab, you can access the System Settings—the master control from which

Figure 6. Kickoff is the new program launcher, a place to find your favorite applications, recently used files, system locations and a whole lot more.

you can change just about anything having to do with your system, from look and feel to sound, printing, networking (including file sharing), display settings and pretty much everything else. It also provides quick access to system locations, including your home folder, the root filesystem, network folders and so on. Recently Used covers both files and applications. The Leave tab is more than merely a way out. You can switch users, log out, shut down or suspend your notebook computer to RAM.

Let's look at another side of running applications. If you know the name of an application, you can, of course, just run it as a command. You also can press Alt-F2, and call up Krunner, a super-powerful program launcher that's a whole lot more (Figure 7). Start typing the name of a command, and Krunner supplies you with options, including program names and icons, before you finish typing. If you see what you like, just click and go.

On Krunner's top right, there's a wrench icon and a small rectangle with a wavy line running through it. Clicking the wrench provides you with a means of configuring Krunner's many plugins. Yes,

Figure 7. Krunner seems to read your mind and offers a number of options as you type your command name.

Krunner can do math, look up recently used Web pages, do a desktop search and run commands too. Click the squiggly line box, and you'll fire up a process monitor (Figure 8). At a glance, you can see every process running under your desktop, including the amount of CPU and memory it is using (both

METRO TORONTO CONVENTION CENTRE, CANADA
CONFERENCE & EXPO : 8 APRIL 2009
MAXIMIZE YOUR BUSINESS SUCCESS
UNIFIED COMMUNICATIONS
VoIP
SECURITY
DATA CENTERS
it360°
Calling all IT professionals and business executives to experience the latest leading industry advancements
www.it360.ca
PRESENTING: ASTERISK AND OPEN TELEPHONY CONFERENCE
IN COOPERATION WITH: CIO Executive Council Leaders Shaping the Future of Business
MEDIA PARTNER: it World Canada IDG CANADA
ASTERISK & OPEN SOURCE TELEPHONY CONFERENCE SPONSOR digium

system and shared). Hover over a program name with your mouse to discover more information. You also can kill a process from the same window.

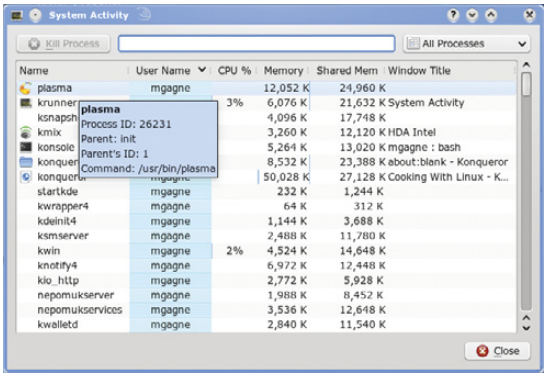

Figure 8. Krunner also provides useful information on your running processes.

If you are the sort of person who needs to build a large list of program favorites, you'll understand the usefulness of what I'm going to show you next. Yes, some of us run an awful lot of applications simultaneously and have, historically, dealt with all those applications by running eight or ten virtual desktops, as opposed to the default four that most distributions give us (some default to two). By default, the top left-hand corner is a hot spot for the mouse. Assuming that you have the desktop effects turned on (true with a compositing capable graphics card), your active applications will all fall flat, side by side, against the desktop background (Figure 9). From there, you easily can see what you have open and quickly switch to it.

Figure 9. Every application on your desktop is just a click away.

Pressing the keyboard shortcut Ctrl-F10 works the same way. You also can choose to view only those applications from your current virtual desktop by pressing Ctrl-F9. Flipping between applications is where you may be used to pressing Alt-Tab. This still works under KDE 4.2, but you may want to select a different behavior and a different effect. I'm rather fond of the cover switch effect where windows flip past you with the current choice facing you (Figure 10). It's a little like flipping through albums (remember vinyl?) or CDs at the music store.

Figure 10. The cover switch effect is one of many ways to flip between applications.

Doing many things at the same time, that old multitasking demon, can create quite the clutter. Imagine you are copying a number of large files from one folder to another or from one system to another. Historically, you would see a number of little progress boxes telling you how each of those copies was progressing. On the right-hand side of the panel, KDE 4.2 now provides an enhanced system tray that multitasks as a notification area, so you can check the progress of those events or just hide them out of the way (Figure 11). The system tray also is configurable with a right-click so you can hide icons you rarely or never use.

Now you've seen the right-hand side of the panel and the left. You might have noticed that on the far left, there's another one of those plasma cashews ready for some action. If you don't see it, it

Figure 11. The notification area can swallow all those progress windows, keeping your desktop clean while you multitask.

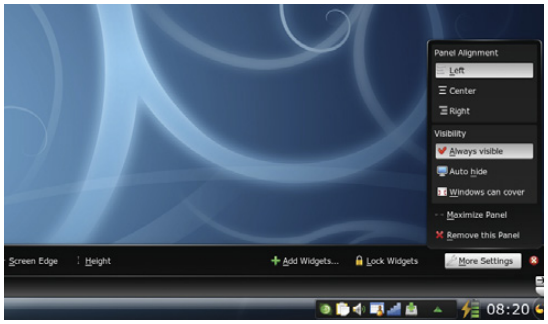

Figure 12. The ultra-configurable panel can be pretty much anything you want.

may be locked. Press Ctrl-L to unlock it and make the cashew appear. Right-click on the cashew, and you can add widgets (plasmoids) to the panel, change basic settings and add or remove another panel. Left-click, and you now are able to change the width and height of the panel. Click the More Settings button, and a host of other settings can be changed—from auto-hiding to alignment to window/panel behavior (Figure 12).

KDE 4.2 is, of course, not only a desktop

Figure 13. Marble, the KDE desktop globe, in now fully integrated with OpenStreetMap.

environment, but also a collection of applications. Some of those applications truly shine in KDE 4.2, providing functionality like never before. One of these applications is Marble, the desktop globe (which I wrote about in the February 2008 issue of

Expert included.

Elizabeth, the product management expert for our rackmount server products, wanted to be here for this picture, but she is **really busy** these days. She's getting ready for the upcoming release of the newest Intel® Xeon® processor technology. Because Silicon Mechanics offers such a comprehensive product line, that means she's readying over 20 different products. Did we mention that she's busy?

Elizabeth knows that features such as Intel QuickPath Architecture, and DDR3 memory support will help our customers get more done in less time using less power than ever before. Want to know how? We'd love to have you call and talk it over with one of the experts at Silicon Mechanics for more information.

When you partner with Silicon Mechanics, you get more than the newest Intel processor technology— you get an expert like Elizabeth.

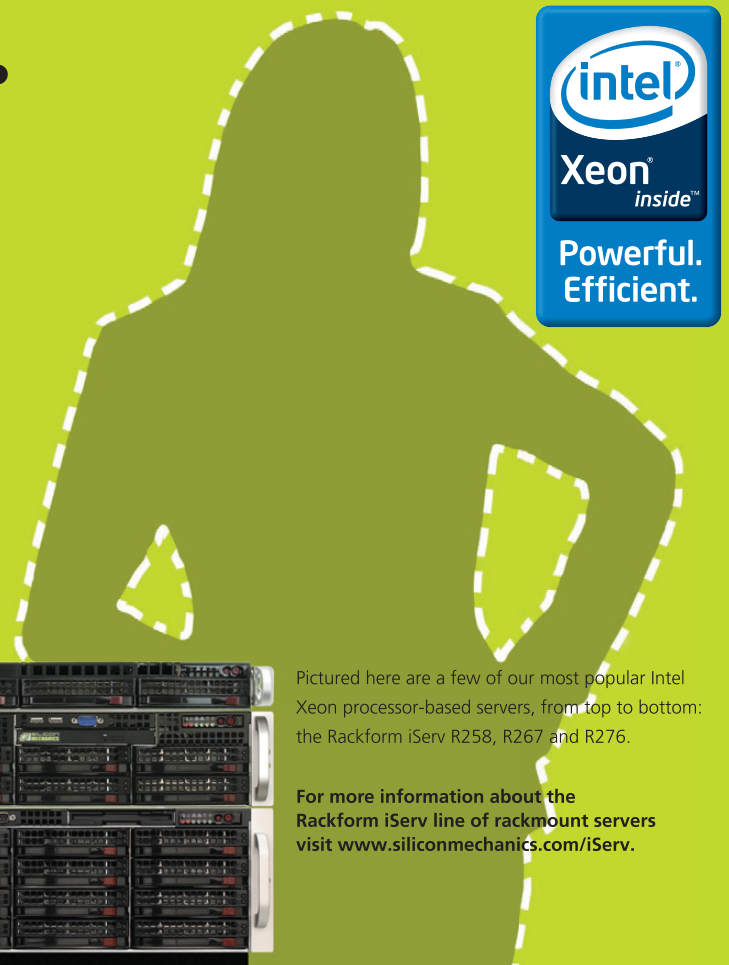

Pictured here are a few of our most popular Intel Xeon processor-based servers, from top to bottom: the Rackform iServ R258, R267 and R276.

For more information about the Rackform iServ line of rackmount servers visit www.siliconmechanics.com/iServ.

SILICON MECHANICS

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

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.

Linux Journal). Marble was more of a toy than anything else in its early days (albeit a very cool toy), but it too has evolved. The current release is fully integrated with the OpenStreetMap Project, which means you can search for a location anywhere on the planet, and if it is in the OpenStreetMap database, you can zoom in on that location (Figure 13).

OpenStreetMap, in case you didn't know, is a collaborative project involving tens of thousands of people worldwide whose goal it is to create and maintain a free editable map of the world. Although the project is considerably more complex than this short explanation can provide, people literally walk the planet with GPS devices, uploading data to the project. There also are organizations that contribute public domain map data, which is then further enhanced by project volunteers. Groups get together for mapping parties. It's great fun. But I digress....

Almost every application has been updated and improved in sometimes subtle, but important ways in KDE 4.2. The task manager in the bottom panel now has a preview of the running applications. Dolphin is the KDE file manager, and although it operates in much the same way that it has for some time, the polish is noticeable there as well. The file view now has a preview of the document or image with a slider along the bottom to increase or decrease the size of the preview (Figure 14).

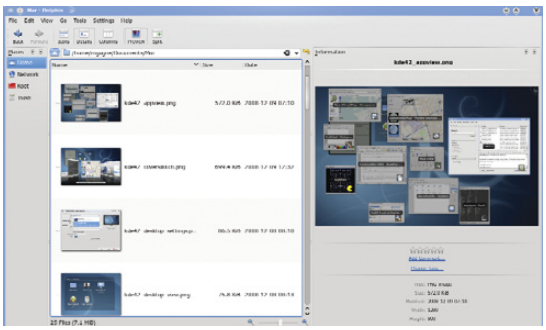

Figure 14. Dolphin now has a cool preview in file mode with a handy slider to define the preview size.

On the admin side, things have improved as well. KDE now comes with PowerDevil, a new power management console that makes mobile life much more flexible (Figure 15). For instance, you can create profiles to define how power management (such as screen blanking, suspend and so on) behaves under different conditions. You might be doing a presentation, for instance, and taking questions while your notebook sits idle. PowerDevil would let you define a profile where the notebook wouldn't just suspend or go dark while it is idle.

There are so many changes in this new KDE that it's difficult to know when to stop. The article

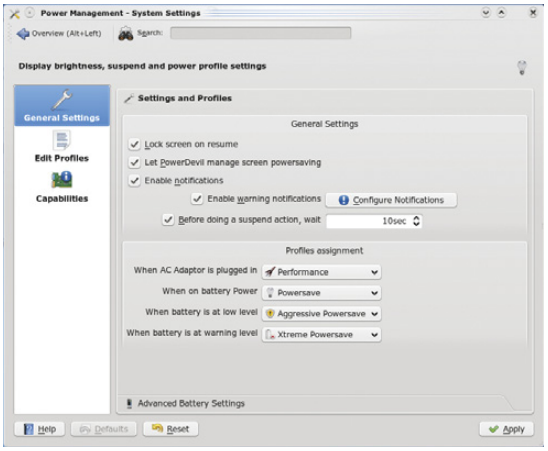

Figure 15. PowerDevil, a new power management system for mobile computing, is powerful and welcome.

clock, however, has other ideas. There are new games, new desktop effects, new configurations and new everything.

It's easy to think of Linux desktop environments as consisting of only GNOME and KDE, but there are many more. Those two, however, GNOME and KDE, certainly are the most popular and powerful. Both environments evolve with each new release. With the release of KDE 4.2, and the 4.X branch in general, the Linux desktop has been dramatically re-imagined, re-engineered and redesigned. Whether you decide it's the environment for you or not, this may well be the pinnacle of Linux desktop evolution—at least for now. Speaking of now, it most certainly is closing time, and I see that François is refilling your glasses a final time. Please, *mes amis*, raise your glasses, and let us all drink to one another's health. *A votre santé! Bon appétit!*■

Marcel Gagné is an award-winning writer living in Waterloo, Ontario. He is the author of the *Moving to Linux* series of books from Addison-Wesley. Marcel is also a pilot, a past Top-40 disc jockey, writes science fiction and fantasy, and folds a mean Origami T-Rex. He can be reached via e-mail at marcel@marcelgagne.com. You can discover lots of other things (including great Wine links) from his Web sites at www.marcelgagne.com and www.cookingwithlinux.com.

Resources

KDE: **www.kde.org**

Marcel's Web Site: **www.marcelgagne.com**

Cooking with Linux: **www.cookingwithlinux.com**

WFTL Bytes!: **wftlbytes.com**

Experience Lightning Without The Thunder℠

WhisperStation-Pro™

1 TERAFLOP IN A COOL, FAST, RELIABLE PLATFORM!

Whether it's Wall Street, Main Street or Your Street, Microway's new Nvidia-powered WhisperStation-Pro is energy-efficient, designed for superior performance, and best of all - QUIET.

Originally designed for a group of power hungry, demanding engineers in the automotive industry, WhisperStation-Pro incorporates two AMD® Opteron™ or Intel® Xeon® quad-core processors and high-efficiency power supplies. Ultra-quiet fans and internal sound-proofing produce a powerful, but silent, computational platform.

WhisperStation-Pro configured with one Quad core processor, 4 GB high speed memory, 250 GB drive, dual-GigE, NVIDIA® Quadro™ FX570 graphics and 20" LCD – **starts at $1995.**

You can have it configured to your exact needs with NVIDIA GeForce® or Quadro graphics adapters (including SLI®), NVIDIA Tesla™ GPU, any Linux distribution, or Windows® dual-boot. Also, there is plenty of room for RAID storage expansion. From a home based workstation for financial wizards, to a superior gaming or design station, WhisperStation-Pro fits the bill and your budget.

Visit www.microway.com for more technical information.

WhisperStation™ 3D Elite SLI For Gamers

Hear Yourself Think Again!

Call our technical sales team at 508-746-7341 and customize your WhisperStation-Pro today.

AMD Opteron 64 · NVIDIA PARTNERFORCE PREMIER · intel Xeon inside™

Microway®

Technology you can count on℠

DAVE TAYLOR

Counting Words and Letters

Count the number of unique words in a document.

I know I have been writing about the basics of working with variables in shell scripts, but I'm going to diverge and address a recent query I received. Okay? (And, hey, write to me.)

"Dear Dave, I seek an edge when I next play *Hangman* or other word games. I want to know what words are most common in the English language and what letters are most common in written material too. If you can show how to do that as a shell script, it'd be useful for your column, but if not, can you point me to an on-line resource? Thanks, Mike R."

Okay, I can tell you up front, Mike, that the secret to playing *Hangman* is to ensure that you have sufficient guesses to get at least 30% of the letters before you're in great peril. Oh, that's not what you seek, is it? The first letter to guess, always, is E, which is the most common letter in the English language. If you have a *Scrabble* set, you also can figure out the frequency of letters, because the points for individual letters are inversely proportional to their frequency. That is, E is worth one point, while the Q and Z—two very uncommon letters in English—are worth ten points each.

If you analyze only *Dracula*, by the way, it turns out that the entire book has only 9,434 unique words.

But, let's work on a shell script to verify and prove all this, shall we?

The first step is to find some written material to analyze. That's easily done by going to one of my favorite places on the Web, the Gutenberg Project. You can pop there too at **www.gutenberg.org**.

With thousands and thousands of books available in free, downloadable form, let's grab only three: *Dracula* by Bram Stoker, *History of the United States* by Charles A. Beard and Mary Ritter Beard, and *Pride and Prejudice* by Jane Austen. They're all obviously a bit older, but that's okay for our purposes. To make life easy, I'll download them as plain text and leave the geeky introduction to the Gutenberg Project at the top of each file too, just for more word variation and, well, because I'm lazy. Okay with you, dear reader?

Here's a quick heads up on the three:

```
$ wc *txt
   16624  163798  874627 dracula.txt
   24398  209289 1332539 history-united-states.txt
   13426  124576  717558 pride-prejudice.txt
   54448  497663 2924724 total
```

Okay, so we have 54,448 lines of text, representing 497,663 words and 2,924,724 characters. That's a lot of text.

Word Frequency

The key to figuring out any of our desired statistics is to realize that the basic strategy we need to use is to break the content down into smaller pieces, sort them, and then use the great `uniq -c` capability, which de-dupes the input stream, counting frequency as it goes. As a shell pipe, we're talking about `sort | uniq -c`, coupled with whatever command we need to break down the individual entities.

For this task, I'm going to use `tr`, like this, to convert spaces to newlines:

```
$ cat *txt | tr ' ' '\
' | head
The
Project
Gutenberg
EBook
of
Dracula,
by
Bram
Stoker
```

Okay, so what happens when we actually unleash the beast on all 54,448 lines of our combined text?

```
$ cat *txt | tr ' ' '\
> ' | wc -l
  526104
```

That's strange. Somehow I would expect that breaking down every line by space delimiter should be fairly close to the word count of `wc`,

but most likely the document has punctuation like "the end. The next" where a double space becomes two, not one line. No worries, though, it'll all vanish once we take the next step.

Now that we have the ability to break down our documents into individual words, let's sort and "uniq" it to see what we see:

```
$ cat *txt | tr ' ' '\
' | sort | uniq | wc -l
   52407
```

But, that's not right. Do you know why?

If you said, "Dude! You need to account for capitalization!", you'd be on the right track. In fact, we need to transliterate everything to lowercase. We also need to strip out all the punctuation as well, because right now it's counting "cat," and "cat" as two different words—not good.

First off, transliteration is best done with a character group rather than with a letter range. In tr, it's a bit funky with the [::] notation:

```
$ echo "Hello" | tr '[:upper:]' '[:lower:]'
hello
```

Stripping out punctuation is a wee bit trickier, but not much. Again, we can use a character class in tr:

```
$ echo "this, and? that! for sure." | tr -d '[:punct:]'
this and that for sure
```

Coolness, eh? I bet you didn't know you could do that! Now, let's put it all together:

```
$ cat *txt | tr ' ' '\
' | tr '[:upper:]' '[:lower:]' |
tr -d '[:punct:]' | sort | uniq | wc -l
   28855
```

So, that chops it down from 52,407 to 28,855—makes sense to me. One more transform is needed though. Let's strip out all lines that don't contain alphabetic characters to eliminate digits. That can be done with a simple `grep -v '[^a-z]'`":

```
$ cat *txt | tr ' ' '\
' | tr '[:upper:]' '[:lower:]' |
tr -d '[:punct:]' | grep -v '[^a-z]' |
sort | uniq | wc -l
   19,820
```

If you analyze only *Dracula*, by the way, it turns out that the entire book has only 9,434 unique words. Useful, eh?

Now, finally, let's tweak things just a bit and see the ten most common words in this corpus:

```
$ cat *txt | tr ' ' '\
' | tr '[:upper:]' '[:lower:]' |
tr -d '[:punct:]' | grep -v '[^a-z]' |
sort | uniq -c | sort -rn | head
29247 the
19925
16995 of
14715 and
13010 to
9293 in
7894 a
6474 i
5724 was
5206 that
```

And, now you know.

Next month, I'll wrap this up by showing how you can analyze individual letter occurrences too, and finally, I'll offer a way to find some great *Hangman* words for stumping your friends.■

Dave Taylor is a 26-year veteran of UNIX, creator of The Elm Mail System, and most recently author of both the best-selling *Wicked Cool Shell Scripts* and *Teach Yourself Unix in 24 Hours*, among his 16 technical books. His main Web site is at www.intuitive.com, and he also offers up tech support at AskDaveTaylor.com. You also can follow Dave on Twitter through twitter.com/DaveTaylor.

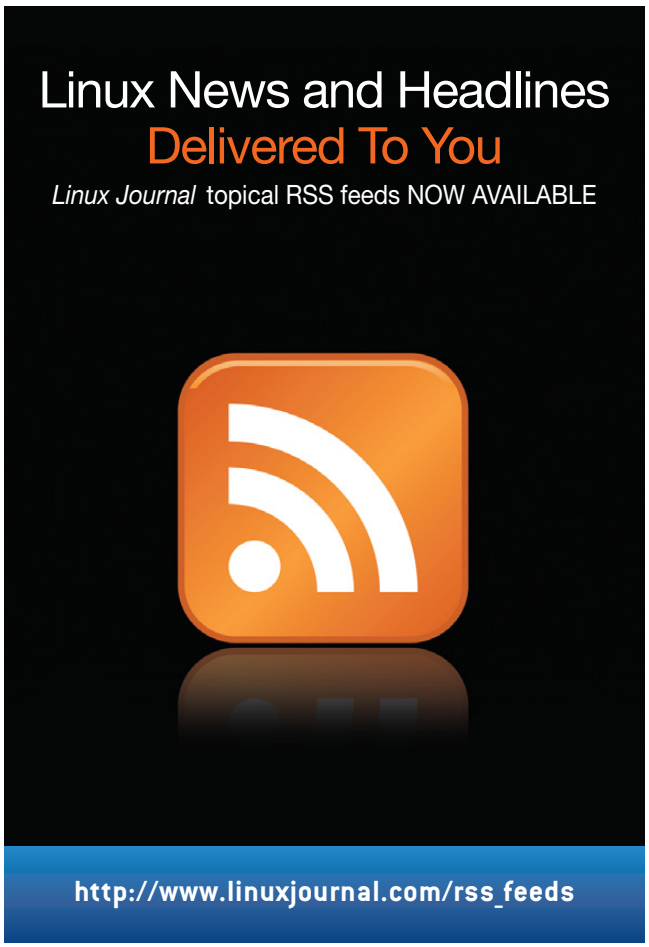

Mental Laziness and Bad Dogma to Avoid

MICK BAUER

Peer pressure is no substitute for common sense.

Gentle readers, I try not to rant at you, really I do. You turn to my column for practical, reliable tips on getting complex security-related tools to work, and judging from the e-mail messages you send me, most of the time I deliver that.

But, I'm human, and now and then, I get really tired of dealing with mental laziness and dogma. It's not because I'm some sort of purist—quite the contrary. Rather, it's because it's *impractical*. Each of us security geeks has a limited amount of energy and political capital, and we can't afford to squander it on positions we can't back up with compelling, plausible risk and threat drivers.

Similarly, although I've got tremendous sympathy for nongeeks who strictly use computers as tools, and who find it (rightly) unreasonable to have to know as much as a system administrator just to be able to print their spreadsheets, Internet use has its price. If you're going to comingle your data with that of practically the entire rest of the world, you need to think about risks now and then, and you need to take the time to learn some simple precautions.

So this month, I need to vent just a little bit about some nagging bits of information security dogma to which security practitioners sometimes cling, and some examples of mental laziness in which end users (especially "power users") sometimes indulge. Your opinions may differ (widely) from mine, and if you take strong exception to any of this, I encourage you to post comments to the Web version of this article or e-mail me directly.

In Defense of Dogma

Before I begin the rant proper, let me acknowledge that to a point, dogma can be useful, in the same way that a parent may now and then find it useful to tell a cantankerous child "the answer is no, because I said so".

Life is short, information security is complicated, and we don't always have the luxury of explaining every rule to every user's satisfaction. Sometimes, it seems to me, it's perfectly appropriate to say, "You can't do that because it violates corporate security policy." The real question is, "Is that a defensible policy?"

So, perhaps my point is not that there is no place in the world for information security dogma, but rather it's that dogma existing only for its own sake is useless. If we can't back up a policy, practice or other security requirement with compelling, risk-based justification, we will fail.

This month's column, therefore, is about some wrong ideas that have somehow ended up being treated as immutable truth among some of my peers, but whose rationales are questionable and tend to cause more harm than good. And, because I don't want anyone to think I'm unduly biased against my colleagues, I'll give equal time to the aforementioned examples of end-user mental laziness as well.

Bad Dogma 1: Changing All Your Passwords Monthly Is Really Important

Consider hapless Hapgood, a typical corporate computer user. At work, Hapgood has to keep track of six different user accounts, each with slightly different password-complexity rules: system A requires a minimum of eight characters containing uppercase and lowercase, punctuation and numbers; system B allows only seven-character passwords, doesn't allow punctuation and so forth.

Due to corporate security policy, within any given 60-day period, Hapgood must change *all six* passwords—a couple of them *twice*. If Hapgood starts choosing passwords that are easy for him to remember but not very hard to guess (for example, his own name with a capital H and zeroes instead of Os), can you really blame him?

I wouldn't. But, which do you suppose is more dangerous: choosing a bad password, or leaving a good password alone for, say, 90 days instead of 30?

Naturally, that depends on what you're worried about. If you're worried about brute-force password attacks in which an attacker cycles through all possible passwords for a given user account, then the more randomized the password, the less likely it will turn up in the password "dictionaries" many attackers employ. In that scenario, short password lifetimes will lower the chance that any given password will be cracked before it expires. But, the password shouldn't be very easily cracked if it's sufficiently complex to begin with. So as it happens, enforcing good password complexity rules is a

better protection against brute-force password attacks.

What if you're worried about Hapgood being fired, but connecting back into the network via a VPN connection and logging back in to his old accounts, in order to exact revenge? Won't a 60-day password lifetime minimize the amount of havoc Hapgood can wreak?

This question is best answered with two other questions. First, why should Hapgood still have access for even *one day* after being fired? Second, if Hapgood's accounts haven't all been de-activated within 60 days, what's to stop him from simply changing his passwords once they expire?

Obviously, in this scenario, password aging is the wrong control on which to fixate. The terminated-employee conundrum can be addressed only by good processes—specifically, the prompt and universal disabling of every terminated employee's account.

There's a third risk people hope will be mitigated by password lifetimes—that a password may be eavesdropped over the network, read off the sticky note attached to someone's monitor or keyboard or otherwise intercepted. This risk is probably more credible than brute-force attacks and user attrition combined.

But even here, if attackers can abuse someone else's access privileges for 29 days without fear of detection, there's probably something seriously wrong with how you're doing things. Furthermore, it may be possible for such attackers to install a keylogger, rootkit or other malware that allows them to intercept the new password, once the intercepted one expires and its rightful owner changes it.

Passwords should, of course, have finite lifetimes. User name/password authentication is a relatively weak form of authentication to begin with, and requiring people to refresh their passwords from time to time certainly makes the attacker's job a little harder. But, compared to password complexity rules and good walkout procedures, password aging achieves less and affects end-user experience more negatively.

Bad Dogma 2: All Digital Certificates Should Expire after One Year

On a related note, consider the digital certificate, which consists of a couple key pairs (one for signing/verifying, another for encrypting/decrypting), identity information (such as your name and organization) and various Certificate Authority signatures. Conventional wisdom says that every digital certificate must have an expiration date, the shorter the better, in case the certificate's owner unexpectedly leaves your organization or the private key is somehow compromised. The consequences of either event could include bogus signatures, illicit logins or worse.

This worst-case scenario assumes two things. First, if the certificate's owner leaves your organization, it may take a while for the certificate to be revoked (and for news of that revocation to propagate to the systems that use certificates). Second, it assumes that the certificate's passphrase can be guessed or brute-force cracked easily.

But, both of these are solvable problems. If you're deploying a Public Key Infrastructure in the first place, you need to configure all systems that use certificates either to download automatically and use Certificate Revocation Lists (CRLs) from your Certificate Authority, or better still, configure them to use the Online Certificate Status Protocol (OCSP). Many events can lead to a certificate's need to be revoked besides reaching some arbitrary expiration date, and managing your certificates diligently and using CRLs or OCSP are the only reliable means of reacting to those events.

Regarding certificate passphrases, setting passphrase complexity requirements is generally no harder for digital certificates than for system passwords. The situation in which it can be most challenging to protect certificate passphrases is when machines use certificates (for example, Web server SSL/TLS certificates), which usually requires either a passphrase-less certificate or a certificate whose passphrase is stored in clear text in some file to which the certificate-using process has read-access privileges.

The bad news is, in that scenario, renewing the server's certificate every year doesn't solve this problem. If it's possible for people to copy a server's certificate once, it's probably possible for people to do so every year, every six months or as often as they need or like. The solution to this problem, rather, is to protect the certificate at the filesystem/OS level, especially its passphrase file, if applicable.

Does that mean certificates *shouldn't* have expiration dates? Of course not! I'm simply saying that, as with password aging, if this is your only protection against user attrition or certificate compromise, you're in big trouble anyhow, so why not employ a *variety* of protections that allow you to relax a little on expiration dates, as you ought to be doing those other things anyhow?

Bad Dogma 3: E-Mail Encryption Is Too Complicated for Ordinary People to Use

For as long as I've worked on information security in large corporations, I've been told that e-mail encryption is only for geeks, and that business users lack the technical skills necessary to cope with it. I've always found this sort of amusing, given that it's usually us geeks who accuse business people of having too-short attention spans.

But, is using PGP or S/MIME really that much more complicated than using, say, LinkedIn? I know which *I* would rather spend time on! (I am, however, an admitted geek.)

How much of the inconvenience in e-mail encryption really falls on end users? Nowadays, I would argue, very little, especially if your organization can support a PGP key server or can incorporate S/MIME certificates into an MS-Exchange Global Address List.

In practice, key management tends to be the biggest headache with e-mail encryption—specifically, getting a valid/current digital certificate or PGP key for each person with which you need to communicate. But, this need not be a big deal if you set things up carefully enough on the back end and give your end users local settings that allow their mail client software to search for, download and update their local copies of other people's keys transparently.

One can go too far, of course, in coddling end users. I've seen organizations issue keys without passphrases, which makes those keys trivially easy to copy and abuse. I've seen other organizations issue passphrase-protected keys, but then send people their new key's initial passphrase via unencrypted e-mail! Obviously, doing things like that can defeat the whole purpose of e-mail encryption.

My point, really, is that modern e-mail encryption tools, which typically support GUI plugins for popular e-mail readers, such as MS Outlook and Squirrelmail, are exponentially simpler to use than the command-line-driven tools of old. Given a modicum of written documentation—a two-page instruction sheet is frequently enough—or even a brief computer-based-training module, nontechnical users can be expected to use e-mail encryption.

This is too valuable a security tool for so much of the world to have given up on!

There, I'm starting to feel better already! But, I'm not done yet. On to some mental laziness that never fails to annoy and frustrate.

Mental Laziness 1: Firewalls Protect You from Your Own Sloppiness

Your DSL router at home has a built-in firewall you've enabled, and your corporate LAN at work has industrial-strength dedicated firewalls. That means, you can visit any Web site or download any program without fear of weirdness, right?

Wrong.

In the age of evil-twin (forged) Web sites, cross-site scripting, spyware and active content, you take a risk every time you visit an untrusted Web site. Your home firewall doesn't know or care *what* your browser pulls, so long as it pulls it via RFC-compliant HTTP or HTTPS. Even Web proxies generally pass the data payloads of HTTP/HTTPS packets verbatim from one session to the other.

This means the site you're visiting may transparently push hostile code at your browser, such as invisible iframe scripts, ActiveX or JavaScript applets (depending on how your browser is configured), or your data may redirected via cross-site scripting and request forgery.

Firewalls are great at restricting traffic by application-protocol type and source and destination IP address, but they aren't great at detecting evil within allowed traffic flows. And nowadays, RFC-compliant HTTP/HTTPS data flows carry everything from the hyptertext "brochureware" for which the Web was originally designed to remote desktop control sessions, full-motion videoconferencing and pretty much anything else you'd care to do over a network.

With or without a firewall, you need to be careful which sites you frequent, which software you install on your system and which information you transmit over the Internet. Just because your nightclub has a bouncer checking IDs at the door doesn't mean you can trust everybody who gets in.

Mental Laziness 2: Firewalls Need to Block Only Inbound Traffic

In olden times, firewalls enforced a very simple trust model: "inside" equals "trusted", and "outside" equals "untrusted". We configured firewalls to block most "inbound" traffic (that is to say, transactions initiated from the untrusted outside) and to allow most "outbound" traffic (transactions initiated from the trusted inside).

Aside from the reality of insider threats, however, this trust model can no longer really be applied to computer systems themselves. Regardless of whether we trust internal *users*, we must acknowledge the likelihood of spyware and malware infections.

Such infections are often difficult to detect (see Mental Laziness 3); and frequently result in infected systems trying to infect other systems, trying to "report for duty" back to an external botnet controller or both.

Suppose users download a new stock-ticker applet for their desktops. But, unbeknownst to them, it serves double duty as a keystroke logger that silently logs and transmits any user names, passwords, credit-card numbers or Social Security numbers it detects being typed on the users' systems and transmits them back out to an Internet Relay Chat server halfway around the world.

Making this scenario work in the attacker's favor depends on several things happening. First, users have to be gullible enough to install the software in the first place, which should be against company

CommunityOne

An open developer conference

march 18–19
marriott marquis hotel
new york

june 1–3
moscone center
san francisco

MySQL™

Linux™

GlassFish™

Apache™

Java™

Cloud

"one heavyweight geek-fest"
—Tim Bray, 2008 speaker

"excellent presenters that know their stuff!"
—2008 attendee

Virtualization

OpenSolaris™

"something for everyone: developers, infrastructure specialists, DBAs, sys admins"
—2008 attendee

Security

CommunityOne is a developer conference focused on open-source innovation and implementation, sponsored by Sun Microsystems. Join fellow developers, technologists and students for a day jam-packed with technical sessions, hands-on demos and a free exchange of ideas. Evaluate dozens of open-source projects enabling the Internet and powering enterprise computing, and advance your skills in technology deep dives.

Register Today!

developers.sun.com/events/communityone
priority code: TUX

Copyright © 2008 Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, CommunityOne, JavaOne, GlassFish, Java, MySQL, NetBeans, OpenSolaris, Sun Studio, and all Java-based marks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

CommunityOne, 999 Skyway Road, Suite 300, San Carlos, CA 94070

Sun microsystems

policy—controlling who installs desktop software and why it is an important security practice. Second, the users' company's firewall or outbound Web proxy has to be not scanning downloads for malicious content (not that it's difficult for an attacker to customize this sort of thing in a way that evades detection).

Finally, the corporate firewall must be configured to allow internal systems to initiate outbound IRC connections. And, this is the easiest of these three assumptions for a company's system administrators and network architects to control.

If you enforce the use of an outbound proxy for all outbound Web traffic, most of the other outbound Internet data flows your users really need probably will be on the back end—SMTP e-mail relaying, DNS and so forth—and, therefore, will amount to a manageably small set of things you need to allow explicitly in your firewall's outbound rule set.

The good news is, even if that *isn't* the case, you may be able to achieve nearly the same thing by deploying personal firewalls on user desktops that allow only outbound Internet access by a finite set of local applications. Anything that end users install without approval (or anything that infects their systems) won't be on the "allowed" list and, therefore, won't be able to connect back out.

Mental Laziness 3: If Your Machine Gets Infected with Malware, You'll Know

Some of us rely on antivirus software less than others. There are good reasons and bad reasons for being more relaxed about this. If you don't use Windows (for which the vast majority of malware is written), if you read all your e-mail in plain text (not HTML or even RTF), if you keep your system meticulously patched, if you disconnect it from the network when you're not using it, if you never double-click e-mail links or attachments, if you minimize the number of new/unfamiliar/untrusted Web sites you visit, and if you install software that comes only from trusted sources, *all* of these factors together may nearly obviate the need for antivirus software.

But, if none of that applies, and you simply assume that in the case of infection, you simply can re-install your OS and get on with your life, thinking you'll notice the infection right away, you're probably asking for trouble.

There was a time when computer crimes were frequently, maybe mostly, motivated by mischief and posturing. Espionage certainly existed, but it was unusual. And, the activities of troublemakers and braggarts tend, by definition, to be very obvious and visible. Viruses, worms and trojans, therefore, tended to be fairly noisy. What fun would there be in infecting people if they didn't know about it?

But, if your goal is not to have misanthropic fun but rather to steal people's money or identity or to distribute spam, stealth is of the essence. Accordingly, the malware on which those two activities depend tends to be as low-profile as possible. A spambot agent will generate network traffic, of course—its job is to relay spam. But, if in doing so it cripples your computer's or your LAN's performance, you'll detect it and remove it all the more quickly, which defeats the purpose.

So, most of us should, in fact, run and maintain antivirus software from a reputable vendor. Antivirus software probably won't detect the activity of malware it didn't prevent infection by—there will always be zero-day malware for which there is no patch or antivirus signature—but it will be infinitely more likely to prevent infection than wishful thinking is.

Conclusion

Thus ends my rant. Now that I've got it out of my system, next month, it's back to more technical stuff. Until then, be safe!■

Mick Bauer (darth.elmo@wiremonkeys.org) is Network Security Architect for one of the US's largest banks. He is the author of the O'Reilly book *Linux Server Security*, 2nd edition (formerly called *Building Secure Servers With Linux*), an occasional presenter at information security conferences and composer of the "Network Engineering Polka".

TECH TIP Watch Live Interrupts

To see the interrupts occurring on your system, run the command:

```
# watch -n1 "cat /proc/interrupts"

           CPU0       CPU1
  0:        330          0   IO-APIC-edge      timer
  1:      11336          0   IO-APIC-edge      i8042
  4:          2          0   IO-APIC-edge
  6:          3          0   IO-APIC-edge      floppy
 ...
NMI:          0          0   Non-maskable interrupts
LOC:    5806923    6239132   Local timer interrupts
 ...
```

The -n1 option passed to watch causes the specified command to be re-run every second.—CHILAKA RAVI KUMAR

The Straight Talk People℠
SINCE 1991

ABERDEEN

VMWARE CERTIFIED SOLUTIONS

ABERDEEN STIRLING 123

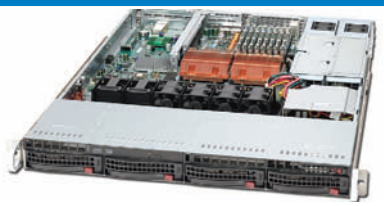

1U Dual Xeon VMware Certified Server
- Up to two Dual-Core or Quad-Core Intel® Xeon® processors
- Up to 32GB 667/533MHz Fully Buffered ECC DDR2 SDRAM
- Up to 4 x Hot-Swap SATA or SAS Hard Drives
- Universal I/O allows for 3 expansion cards in 1U
- Pre-installed VMware® ESXi on Disk-on-Module
- 650W High-efficiency Redundant Power Supply
- **5-Year Warranty**

Starting at **$1,999**

ABERDEEN STIRLING 229

2U Dual Xeon VMware Certified Server
- Up to two Dual-Core or Quad-Core Intel Xeon processors
- Up to 128GB 800/667/533MHz Fully Buffered ECC DDR2 SDRAM
- Up to 8 x 1TB (8.0TB) Hot-Swap SATA Hard Drives
- Up to 7 x Low-Profile Expansion Slots
- Pre-installed VMware ESXi on Disk-on-Module
- 700W High-efficiency Redundant Power Supply
- **5-Year Warranty**

Starting at **$2,425**

ABERDEEN STIRLING 428T

4U Dual Xeon VMware Certified Server
- Up to two Dual-Core or Quad-Core Intel Xeon processors
- Up to 64GB 667/533MHz Fully Buffered ECC DDR2 SDRAM
- Up to 8 x 1TB (8.0TB) Hot-Swap SATA Hard Drives
- Up to 6 x Full Height Expansion Slots
- Pre-installed VMware ESXi on Disk-on-Module
- 800W High-efficiency Redundant Power Supply
- **5-Year Warranty**

Starting at **$2,350**

ABERDEEN STIRLING 132T

1U Twin Node VMware Certified Server
- Up to two Dual-Core or Quad-Core Intel Xeon processors/node
- Twin Nodes allows for up to 4 processors & 16 cores in 1U
- Up to 64GB 800/667/533MHz Fully Buffered ECC DDR2/node
- Up to 2 x 1TB Hot-Swap SATA Hard Drives per node
- Pre-installed VMware ESXi on Disk-on-Module
- 980W High-efficiency Power Supply
- **5-Year Warranty**

Starting at **$2,880**

ABERDEEN STIRLING 244

2U Quad Xeon MP VMware Certified Server
- Up to four Quad-Core or Six-Core Intel Xeon MP processors
- Quad Six-Core allows for 24 processor cores in 2U
- Up to 192GB 667/533MHz Fully Buffered ECC DDR2 SDRAM
- Up to 6 x Hot-Swap SATA or SAS Hard Drives
- Pre-installed VMware ESXi on Disk-on-Module
- 1200W High-efficiency Redundant Power Supply
- **5-Year Warranty**

Starting at **$5,615**

ABERDEEN STIRLING 444

4U Quad Xeon MP VMware Certified Server
- Up to four Quad-Core or Six-Core Intel Xeon MP processors
- Quad Six-Core allows for 24 processor cores in 4U
- Up to 192GB 667/533MHz Fully Buffered ECC DDR2 SDRAM
- Up to 5 x Hot-Swap SATA or SAS Hard Drives
- Pre-installed VMware ESXi on Disk-on-Module
- 1200W High-efficiency Redundant Power Supply
- **5-Year Warranty**

Starting at **$6,625**

ABERDEEN ISCSI XDAS

DAS VMware Certified Expandable Storage
- IP SAN Solution
- Single or Redundant Controller
- Expandable up to 64TB in a single array
- 2U/12 Bay and 3U/16 Bay Models available
- SAS or SATA Hard Drive Support
- Fault-tolerant Modular Hardware Design
- **5-Year Warranty**

Starting at **$7,995**

ABERDEEN FIBRE XDAS

DAS VMware Certified Expandable Storage
- Hardware RAID5 and RAID6 engine by dedicated ASIC400
- Single or Redundant Controller
- Expandable up to 64TB in a single array
- 2U/12 Bay and 3U/16 Bay Models available
- SAS or SATA Hard Drive Support
- Fault-tolerant Modular Hardware Design
- **5-Year Warranty**

Starting at **$8,495**

Intel, Intel Logo, Intel Inside, Intel Inside Logo, Pentium, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. For terms and conditions, please see www.aberdeeninc.com/abpoly/abterms.htm. lj029

888-297-7409
www.aberdeeninc.com/lj029

KYLE RANKIN

When Disaster Strikes: Hard Drive Crashes

All is not necessarily lost when your hard drive starts the click of death. Learn how to create a rescue image of a failing drive while it still has some life left in it.

The following is the beginning of a series of columns on Linux disasters and how to recover from them, inspired in part by a Halloween *Linux Journal* Live episode titled "Horror Stories". You can watch the original episode at **www.linuxjournal.com/video/linux-journal-live-horror-stories**.

Nothing teaches you about Linux like a good disaster. Whether it's a hard drive crash, a wayward `rm -rf` command or fdisk mistakes, there are any number of ways your normal day as a Linux user can turn into a nightmare. Now, with that nightmare comes great opportunity: I've learned more about how Linux works by accidentally breaking it and then having to fix it again, than I ever have learned when everything was running smoothly. Believe me when I say that the following series of articles on system recovery is hard-earned knowledge.

Treated well, computer equipment is pretty reliable. Although I've experienced failures in just about every major computer part over the years, the fact is, I've had more computers outlast their usefulness than not. That being said, there's one computer component you can almost count on to fail at some point—the hard drive. You can blame it on the fast-moving parts, the vibration and heat inside a computer system or even a mistake on a forklift at the factory, but when your hard drive fails prematurely, no five-year warranty is going to make you feel better about all that lost data you forgot to back up.

The most important thing you can do to protect yourself from a hard drive crash (or really most Linux disasters) is back up your data. Back up your data! Not even a good RAID system can protect you from all hard drive failures (plus RAID doesn't protect you if you delete a file accidentally), so if the data is important, be sure to back it up. Testing your backups is just as important as backing up in the first place. You have not truly backed up anything if you haven't tested restoring the backup. The methods I list below for recovering data from a crashed hard drive are much more time consuming than restoring from a backup, so if at all possible, back up your data.

Now that I'm done with my lecture, let's assume that for some reason, one of your hard drives crashed and you did not have a backup. All is not necessarily lost. There are many different kinds of hard drive failure. Now, in a true hard drive crash, the head of the hard drive actually will crash into the platter as it spins at high speed. I've seen platters after a head crash that are translucent in sections as the head scraped off all of the magnetic coating. If this has happened to you, no command I list here will help you. Your only recourse will be one of the forensics firms out there that specialize in hard drive recovery. When most people say their hard drive has crashed, they are talking about a less extreme failure. Often, what has happened is that the hard drive has developed a number of bad blocks—so many that you cannot mount the filesystem—or in other cases, there is some different failure that results in I/O errors when you try to read from the hard drive. In many of these circumstances, you can recover at least some, if not most, of the data. I've been able to recover data from drives that sounded *horrible* and other people had completely written off, and it took only a few commands and a little patience.

Create a Recovery Image

Hard drive recovery works on the assumption that not all of the data on the drive is bad. Generally speaking, if you have bad blocks on a hard drive, they often are clustered together. The rest of the data on the drive could be fine if you could only access it. When hard drives start to die, they often do it in phases, so you want to recover as much data as quickly as possible. If a hard drive has I/O errors, you sometimes can damage the data further if you run filesystem checks or other repairs on the device itself. Instead, what you want to do is create a complete image of the drive, stored on good media, and then work with that image.

A number of imaging tools are available for Linux—from the classic dd program to advanced GUI tools—but the problem with most of them is that they are designed to image healthy drives. The

problem with unhealthy drives is that when you attempt to read from a bad block, you will get an I/O error, and most standard imaging tools will fail in some way when they get an error. Although you can tell dd to ignore errors, it happily will skip to the next block and write nothing for the block it can't read, so you can end up with an image that's smaller than your drive. When you image an unhealthy drive, you want a tool designed for the job. For Linux, that tool is ddrescue.

ddrescue or dd_rescue

To make things a little confusing, there are two similar tools with almost identical names. dd_rescue (with an underscore) is an older rescue tool that still does the job, but it works in a fairly basic manner. It starts at the beginning of the drive, and when it encounters errors, it retries a number of times and then moves to the next block. Eventually (usually after a few days), it reaches the end of the drive. Often bad blocks are clustered together, and in the case when all of the bad blocks are near the beginning of the drive, you could waste a lot of time trying to read them instead of recovering all of the good blocks.

The ddrescue tool (no underscore) is part of the GNU Project and takes the basic algorithm of dd_rescue further. ddrescue tries to recover all of the good data from the device first and then divides and conquers the remaining bad blocks until it has tried to recover the entire drive. Another added feature of ddrescue is that it optionally can maintain a log file of what it already has recovered, so you can stop the program and then resume later right where you left off. This is useful when you believe ddrescue has recovered the bulk of the good data. You can stop the program and make a copy of the mostly complete image, so you can attempt to repair it, and then start ddrescue again to complete the image.

Prepare to Image

The first thing you will need when creating an image of your failed drive is another drive of equal or greater size to store the image. If you plan to use the second drive as a replacement, you probably will

MAYBE ONE DAY SECURITY WILL BE EASY.
UNTIL THEN, THERE'S BLACK HAT.

Black Hat is a community that brings together security professionals and visionaries from everywhere: government, academia, corporate security and the underground. We exist to promote and improve information security by facilitating the timely sharing of accurate, actionable security knowledge.

Join us for a full slate of hands-on training sessions and four exciting tracks of presentations from the best minds in the information security space, all with a special focus on the needs of government and infrastructure security.

Black Hat
BRIEFINGS AND TRAINING

BLACK HAT DC 2009
FEBRUARY 16-19
ARLINGTON, VIRGINIA

REGISTER NOW AT BLACKHAT.COM
USE CODE BHDC9LJMAG TO SAVE $100 ON BRIEFINGS REGISTRATION

Platinum Sponsor
NORTHROP GRUMMAN
nitrosecurity

Gold Sponsors
BlackBerry
NORMAN
GIGAMON
SAINT
IOActive
splunk>
Looking Glass SYSTEMS
StillSecure
Microsoft
TRUSTED
NETWITNESS

want to image directly from one device to the next. However, if you just want to mount the image and recover particular files, or want to store the image on an already-formatted partition or want to recover from another computer, you likely will create the image as a file. If you do want to image to a file, your job will be simpler if you image one partition from the drive at a time. That way, it will be easier to mount and fsck the image later.

The ddrescue program is available as a package (ddrescue in Debian and Ubuntu), or you can download and install it from the project page. Note that if you are trying to recover the main disk of a system, you clearly will need to recover either using a second system or find a rescue disk that has ddrescue or can install it live (Knoppix fits the bill, for instance).

As you can see, a great thing about ddrescue is that it gives you constantly updating output, so you can gauge your progress as you rescue the partition.

Run ddrescue

Once ddrescue is installed, it is relatively simple to run. The first argument is the device you want to image. The second argument is the device or file to which you want to image. The optional third argument is the path to a log file ddrescue can maintain so that it can resume. For our example, let's say I have a failing hard drive at /dev/sda and have mounted a large partition to store the image at /mnt/recovery/. I would run the following command to rescue the first partition on /dev/sda:

```
$ sudo ddrescue /dev/sda1 /mnt/recovery/sda1_image.img
/mnt/recovery/logfile
Press Ctrl-C to interrupt
Initial status (read from logfile)
rescued:        0 B,  errsize:    0 B,  errors:       0
Current status
rescued:   349372 kB,  errsize:    0 B,  current rate: 19398 kB/s
   ipos:   349372 kB,   errors:    0,    average rate: 16162 kB/s
   opos:   349372 kB
```

Note that you need to run ddrescue with root privileges. Also notice that I specified /dev/sda1 as the source device, as I wanted to image to a file. If I were going to output to another hard drive device (like /dev/sdb), I would have specified /dev/sda instead. If there were more than one partition on this drive that I wanted to recover, I would repeat this command for each partition and save each as its own image.

As you can see, a great thing about ddrescue is that it gives you constantly updating output, so you can gauge your progress as you rescue the partition. In fact, in some circumstances, I prefer using ddrescue over dd for regular imaging as well, just for the progress output. Having constant progress output additionally is useful when considering how long it can take to rescue a failing drive. In some circumstances, it even can take a few days, depending on the size of the drive, so it's good to know how far along you are.

Repair the Image Filesystem

Once you have a complete image of your drive or partition, the next step is to repair the filesystem. Presumably, there were bad blocks and areas that ddrescue could not recover, so the goal here is to attempt to repair enough of the filesystem so you at least can mount it. Now, if you had imaged to another hard drive, you would run the fsck against individual partitions on the drive. In my case, I created an image file, so I can run fsck directly against the file:

```
$ sudo fsck -y /mnt/recovery/sda1_image.img
```

I'm assuming I will encounter errors on the filesystem, so I added the -y option, which will make fsck go ahead and attempt to repair all of the errors without prompting me.

Mount the Image

Once the fsck has completed, I can attempt to mount the filesystem and recover my important files. If you imaged to a complete hard drive and want to try to boot from it, after you fsck each partition, you would try to mount them individually and see whether you can read from them, and then swap the drive into your original computer and try to boot from it. In my example here, I just want to try to recover some important files from this image, so I would mount the image file loopback:

```
$ sudo mount -o loop /mnt/recovery/sda1_image.img /mnt/image
```

Now I can browse through /mnt/image and hope that my important files weren't among the corrupted blocks.

Method of Last Resort

Unfortunately in some cases, a hard drive has far too many errors for fsck to correct. In these situations, you might not even be able to mount the filesystem at all. If this happens, you aren't

necessarily completely out of luck. Depending on what type of files you want to recover, you may be able to pull the information you need directly from the image. If, for instance, you have a critical term paper or other document you need to retrieve from the machine, simply run the strings command on the image and output to a second file:

```
$ sudo strings /mnt/recovery/sda1_image.img >
/mnt/recovery/sda1_strings.txt
```

The sda1_strings.txt file will contain all of the text from the image (which might turn out to be a lot of data) from man page entries to config files to output within program binaries. It's a lot of data to sift through, but if you know a keyword in your term paper, you can open up this text file in `less`, and then press the / key and type your keyword in to see whether it can be found. Alternatively, you can grep through the strings file for your keyword and the surrounding lines. For instance, if you were writing a term paper on dolphins, you could run:

```
$ sudo grep -C 1000 dolphin /mnt/recovery/sda1_strings.txt >
/mnt/recovery/dolphin_paper.txt
```

This would not only pull out any lines containing the word dolphin, it also would pull out the surrounding 1,000 lines. Then, you can just browse through the dolphin_paper.txt file and remove lines that aren't part of your paper. You might need to tweak the -C argument in grep so that it grabs even more lines.

In conclusion, when your hard drive starts to make funny noises and won't mount, it isn't necessarily the end of the world. Although ddrescue is no replacement for a good, tested backup, it still can save the day when disaster strikes your hard drive. Also note that ddrescue will work on just about any device, so you can use it to attempt recovery on those scratched CD-ROM discs too.■

Kyle Rankin is a Senior Systems Administrator in the San Francisco Bay Area and the author of a number of books, including *Knoppix Hacks* and *Ubuntu Hacks* for O'Reilly Media. He is currently the president of the North Bay Linux Users' Group.

SMALL, EFFICIENT COMPUTERS WITH PRE-INSTALLED UBUNTU.

GS-Lo8 Fanless Pico-ITX System
Ultra-Compact, Full-Featured Computer
Excellent for Industrial Applications

3677 Intel Core 2 Duo Mobile System
Range of Intel-Based Mainboards Available
Excellent for Mobile & Desktop Computing

DISCOVER THE ADVANTAGE OF MINI-ITX.
Selecting a complete, dedicated platform from us is simple: Pre-configured systems perfect for both business & desktop use, Linux development services, and a wealth of online resources.

ubuntu solution provider

LOGIC SUPPLY
www.logicsupply.com

Green Rack's Eco Enclosures

Going green is becoming a well-worn path to saving green—greenbacks, that is. In that spirit, green-IT specialist Green Rack Systems has announced Eco Enclosures, a new line of data-center equipment designed to reduce IT budgets while reducing environmental impact. The turnkey Eco Enclosures solutions, which can be customized to customer needs of any size, contain chassis that are made from recycled materials and are fully recyclable. To save power, Eco Enclosures contain features such as "low-wattage multicore CPUs, low-voltage memory, low-power hard drives and ultra-low consumption power supplies". Green Rack also offers solar-powered Web hosting and custom-built data centers upon request.

www.greenracksystems.com

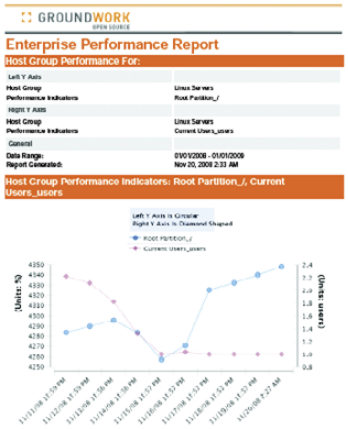

GroundWork Open Source's GroundWork Monitor

GroundWork Monitor is a free, open-source, Nagios-based VMware virtual appliance for network management. It runs on CentOS and ships with a wide range of additional network tools. The new version 5.3 adds features such as updated Eclipse BIRT reporting views, Nagios 3.0.5 and the ability to monitor more than 1,000 different types of devices; automatic notifications of patches, updates and news; voluntary sharing of usage statistics; autodiscovery functionality and increased scalability. The Community Edition is available for download from GroundWork's Web site.

www.groundworkopensource.com

TRENDnet's TEW-635BRM NADSL 2/2+ Modem Router

Ditch one ugly black box by deploying TRENDnet's new TEW-635BRM NADSL 2/2+ Modem Router. The wireless-n device combines modem and router functions and operates at 300Mbps, or 12x the speed and 4x the coverage of wireless g. The switch includes four ports. Key features of the TEW-635BRM include one-touch synchronized Wi-Fi Protected Setup (WPS), the latest in wireless encryption, a double firewall and advanced antenna technology (MIMO) that delivers broad coverage and minimizes dead spots. The device also sports green credentials, such as ENERGY STAR and RoHS compliance—the former related to power consumption and the latter to restricting hazardous substances in electronic equipment.

www.trendnet.com

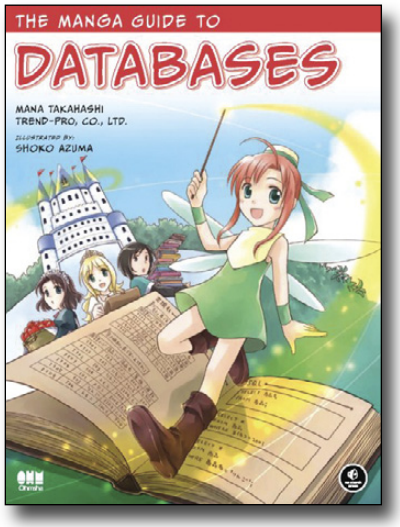

Mana Takahashi and Shoko Azuma's *The Manga Guide to Databases* (No Starch)

Learning databases can be fun? Such is the premise of Mana Takahashi and Shoko Azuma's new book *The Manga Guide to Databases*, published by No Starch Press. The book uniquely fuses Japanese-style comics with serious educational content on databases. It tells the story of Princess Ruruna's challenges in managing the Kingdom of Kod's humongous fruit-selling empire. Tico the fairy teaches the princess how to simplify her data management, and together they design a relational database. They cover concepts such as the the entity-relationship model, basic database operations, SQL statements, database tuning, security, concurrency and replication. Other features include examples and exercises (with answer keys) and an appendix of frequently used SQL statements.

www.nostarch.com

Wesley Chun's *Python Fundamentals Video Training* (Prentice Hall)

If you want to learn Python and your learning style is visual, offer up a gleeful shout for Wesley Chun's new instructional video *Python Fundamentals*. Part of Prentice Hall's LiveLessons series, the video course mirrors topics covered in its sister publication, *Core Python Programming*, Second Edition. Covered are both Python fundamentals, such as syntax basics and standard types and operations, and advanced topics, such as Python's memory model and object-oriented programming. The publisher states that both new and experienced users will find the product useful. The companion booklet contains review questions and coding exercises.

www.informit.com

Round Solutions' AarLogic C10/3

Hot off the assembly line at Round Solutions is the new AarLogic C10/3, a diminutive Linux-based breadboard covering a mere 104mm x 63mm of real estate. The board sports quadband GPRS and SiRF3 GPS modules; interfaces for USB, RS-232 and Ethernet; SD-card reader; 4MB of RAM and dual ARM processors. The processor module is responsible for GSM functionality and for applications executable under embedded Linux. Connectivity options via the board's 192-pin socket include not only keyboards, digital cameras and reading devices but also WLAN, Bluetooth and GPS components. Round Solutions adds that the AarLogic C10/3 can be easily coupled with a camera or environmental sensors. Ideal applications include self-sufficient positioning and monitoring systems due to the wireless provision of spatial coordinates.

www.roundsolutions.com

Opengear's EMD 5000

From the "ounce of prevention" department comes Opengear's EMD 5000, a new environmental management solution for continuous monitoring of environmental conditions at the rack level. The EMD 5000 protects critical data-center assets from hazards, such as heat, humidity, smoke, water leaks, cabinet intrusion and other environmental conditions, by providing facilities managers with extensive logging and alert facilities. Users can view ambient temperature and humidity of the remote environment and set the device to send progressive alarms automatically from warning levels to critical alerts. The EMD also monitors the status of two external dry contacts that can be connected to a smoke detector, water detector, vibration or open-door sensor.

www.opengear.com

IBM, Virtual Bridges and Canonical's Virtual Desktop Solution

The trio of IBM, Virtual Bridges and Canonical have joined forces to offer a Linux-based virtual PC solution that the firms claim halves the costs of licensing and maintaining a desktop. The solution hosts multiple virtual Linux desktops on a server, replacing Microsoft Windows with Ubuntu Linux and Microsoft Office with open-standards-based alternatives, such as IBM Lotus Symphony. The virtual desktop, called Virtual Enterprise Remote Desktop Environment (VERDE), is provided by Virtual Bridges. IBM cites reduced TOC related to licensing, hardware upgrades, power consumption and support requirements. The solution is a key component in IBM's offerings for entities in financial services and in the public sector.

www.ibm.com/linux, www.canonical.com, www.vbridges.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

Warsow—Fast-Paced Cartoon Combat

(www.warsow.net)

In the multitude of first-person shooter projects available comes a particularly solid one: *Warsow*. *Warsow* is based on the Qfusion 3-D engine (itself a modification of the *Quake 2* GPL engine); however, it runs as a completely standalone package and has a solid feel, avoiding the tackier drawbacks of a simple mod. For those of you about to say, "I've seen it all before", hold on, because this is a particularly solid outing with some fresh approaches to game dynamics in an already-stale genre.

Warsow

Warsow has two elements in particular that make it stand out from the rest: speed and motion. *Warsow* is all about how you move around in the 3-D world. It's about fluidity, motion and some interesting changes to game balance. Particularly interesting are moves, such as the "Wall Jump", where pressing a special key when touching a wall allows you to rebound while jumping, or probably the game's main dynamic, "Bunny Hopping". Bunny Hopping has been in first-person games since you could first jump, but *Warsow* adds the element of increased momentum and speed, allowing for a slew of new gameplay tactics and design elements.

Don't be put off by seeing "*Quake 2*" either, Qfusion is not an old and ugly engine destined to turn out some clunky old game that looks blockier than a LEGO factory with clumsy control. *Warsow* is an elegant title including great architecture, gameplay and feel, with its own unique cards brought to the table. *Warsow* has a unique approach to weapons with two types of ammo: the stock ammo that comes with a weapon (weaker) and stronger ammo once more is collected. Aesthetics also play a large part, in particular, a cell shading look similar to manga and the like, lending the game a feel of something like a cross between *Quake III* and *Nerf Arena Blast*. Part of this cell shading ideal is to remove the ultra-realistic, gritty feel of most modern shooters and to reduce the violent content and feel with something more lighthearted with a comic inspiration (which is, indeed, a welcome relief).

Despite the old *Quake 2* base, Qfusion's modifications allow for some great architecture, as seen in *Warsow*.

Cell shading makes *Warsow*'s environment much more colorful and lighthearted than today's standard grizzly fare of action games.

Installation Installing *Warsow* is very easy. Available on the Web site is a unified package containing both Windows and Linux binaries. Download this, and extract it to somewhere convenient. Open the new folder either with a file manager or a terminal if you're the minimalist type. The solitary inconvenience in this package is that you'll have to flag two files as executable: the warsow script and the platform binary that suits your system. For Linux users with an Intel-based machine, you have the choice of warsow.i386 or warsow.x86_64, for 32- and 64-bit systems, respectively.

If you're using a file manager (I use Konqueror for this example, other file managers should be similar), right-clicking on the script and the binary, and choosing Properties and then the Permissions tab will show you the options you need. Check the box for Is executable, and you should be ready to go simply by left-clicking on the warsow script when you're done. For those using a terminal, this should do the trick:

```
$ chmod u+x warsow warsow.i386
```

Once done, start the game by entering:

```
$ warsow
```

Usage First things first, I'm afraid that *Warsow* is a multiplayer-only affair—sorry. However, for those looking to refine their skills without other humans, in-game bots are available (see the game's documentation for more details). Before you jump head on into the action, check out the available tutorials. These are clever presentations using the game itself, but instead of you being in control, it puts the movement "on rails" so to speak, and a voice-over guides you through what is happening.

Once you're confident enough to start the game itself, the controls are the standard FPS affair with WADS controlling the movement, and the Spacebar for jumping, steering and looking around. Shooting is done with the mouse, as well as with the "Special" button, which is used for

dashing, wall jumping and the like. All of the controls are re-assignable, however, and it's well worth customizing it to your own needs as well as checking out the game's other controls.

When you're ready, choose join game to search for an arena to play in, or alternatively, you can host one yourself. At first, join game probably will come up with nothing, so you will have to click search down at the bottom to browse for new games. Choose the server that sounds best for you (look for one with other players if you can, obviously), and if you don't have the map installed, *Warsow* will download it from that game's server.

At the time of this writing, *Warsow* is at 0.40 status, yet the gameplay is seriously solid. There are a few problems here and there, such as the occasional menu quirk and jolts with the sound, but the level of problems in the game are normally what you'd associate with something close to full release instead of an early demonstration. I imagine that *Warsow* probably will add things like single-player skirmishes before it gets to something like 0.9 status, but it's already a fantastic piece of work even for the fussiest of players. Keep an eye on this one, and any programming houses, keep an eye on these coders!

libdmtx—Data Matrix Barcode Scanning

(www.libdmtx.org)

I realise I tend to cover wacky things like molecule imaging, telekinesis and 3-D knitting software, but this is something that actually may be of genuine industrial use in everyday life. libdmtx is an open-source project dedicated to providing tools for reading and writing 2-D Data Matrix barcodes. The Data Matrix standard (**en.wikipedia.org/wiki/Data_Matrix**) is gaining widespread popularity due to its impressive features, but it may be of particular interest to the FOSS community because it's unencumbered by patents and royalty-free (thus, free to use and distribute). Also, the existing proprietary solutions can be quite expensive, and libdmtx now has reached a point where it realistically can save some users six-digit savings

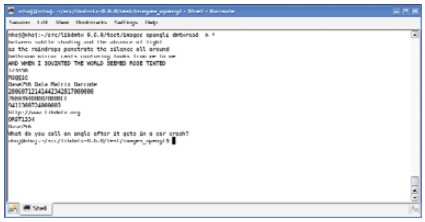

Data Matrix barcodes also can hold secret messages along with the usual barcode data as shown here.

every year.

Data Matrix barcodes have been around since the 1980s, but for years, they were used only to mark electronic components. More recently, they have been adopted by a wide variety of industries in the US and Europe, and they are becoming especially popular with mobile phone developers due to their affinity to work with small digital cameras. Most US readers instantly will recognize Data Matrix barcodes, as they appear on most first-class mail delivered by the US Postal Service. Curious readers can snap a photo of their mail with a camera or Webcam and scan it with libdmtx without purchasing any special hardware (it also works well with faxed and scanned images).

Installation Installing libdmtx is fairly straightforward with either a Debian package available under the name of libdmtx-utils or a source tarball. For those installing via source, compiling is basically the standard affair of:

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

And, as root or sudo:

```
# make install
```

However, the configure script did come up with a dependency you probably won't have installed by default, GraphicsMagick. GraphicsMagick is in many distro repositories though, and to get past the configure script, I had to install libgraphicsmagick1 and libgraphicsmagick1-dev from the Ubuntu archive.

Once you have libdmtx compiled, before you can run the program, you probably will need to run the following command (as root or sudo):

```
# ldconfig
```

libdmtx has the ability to find and decrypt barcodes under a variety of trying conditions.

Usage I cover only very basic usage in reading barcodes for now, but libdmtx also will write barcodes along with a bunch of other features that make it worth checking the man pages. First, grab an image to test. If you have a photo of a barcode around, great stuff, use that. Otherwise, some test images are available from the source tarball under the folder test/images_opengl, which cover a variety of different situations and tricky tests on libdmtx's abilities. Once you're ready to go, use the following command:

```
$ dmtxread nameofimage.png
```

And, that's pretty much all you need to do. dmtxread will spend a few seconds analyzing the image you've given it, and if it finds a matrix barcode, it then outputs the contained text to the terminal. Check the screenshot for some of the hidden messages and real-world codes that you can contain within a barcode.

What really intrigued me about this project is that you can recover barcode data from old pictures that never would have been meant for the purpose originally. And, the James Bond in me gets a kick out of knowing you can hide a message in a barcode in a seemingly unrelated picture as a covert method of communication—neat! Although this has just a command-line utility for now, it's really only a basic program on top of a very clever and versatile library. This project is begging for a GUI front end, at which point, it could make some serious inroads and savings in the real industrial world.

whohas—Package Finder

(www.philippwesche.org/200811/whohas/intro.html)

Finally, we have a tool that will end some serious headaches, whohas. According to the project's readme file:

> whohas is a command-line tool that allows querying several package lists at once—currently supported are Arch, Debian, Gentoo and Slackware. whohas is written in Perl and was designed to help package maintainers find ebuilds, pkgbuilds and similar package definitions from other distributions to learn from. However, it also can be used by normal users who want to know: what distribution has packages available for apps upon which the user depends and what version of a given package is in use in each distribution or in each release of a distribution (implemented only for Debian).

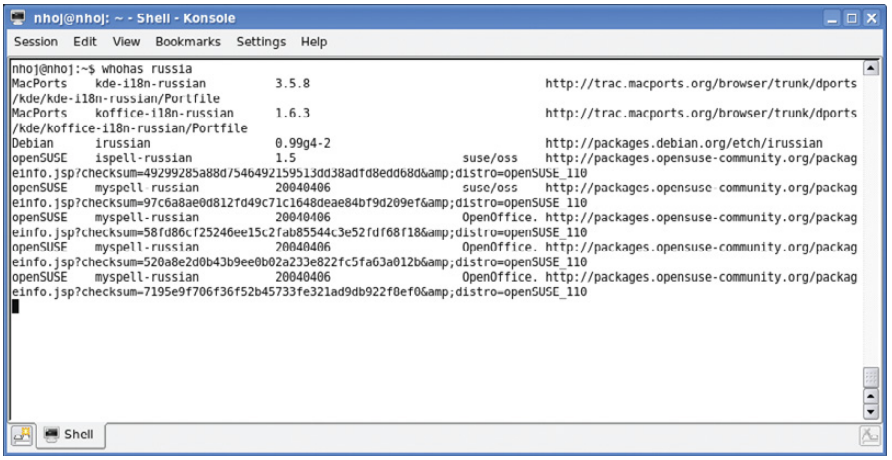

whohas makes finding obscure packages a breeze.

Installation whohas is genuinely very easy to install and is unlikely to have unresolved dependencies on most systems. Compilation isn't a worry either, as whohas is merely a Perl script that gets copied into your /usr/bin directory. To install whohas, grab the latest tarball from the whohas Web site, extract the contents, and open a terminal in the new folder. Then, as root or sudo, enter the following command:

```
# ./install.sh
```

Usage Once whohas is installed, using the program is as easy as entering:

```
$ whohas nameofpackage
```

You needn't be super-specific either. A simple search like "quake", "audacity" or "chartr" will do fine—no need for entering something like "chartr_0.16_i386.deb". whohas then scans a number of repositories and prints the results out to screen one by one. If nothing comes up at first, don't despair; it still might be searching. whohas also provides URLs to more details about the package, so project maintainer Philipp Wesche recommends using a terminal that recognises hyperlinks and allows easy forwarding to the browser. And, as with most *nix command-line programs of this nature, the results can be piped through to grep and the like for further refinement (and if you know what that means, you don't need me to explain how it works).

I'm hoping whohas will become an everyday tool the way we use something like whereis or grep. whohas is a fantastic little project that should have existed years ago, and hopefully, it will make its way into many distributions by default.■

John Knight is a 24-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.

On the Web, Articles Talk!

Dual booting is a sometimes necessary evil and very inconvenient. What if you could run your Windows partition in a virtual machine, so you wouldn't have to worry about rebooting anymore? With VMware Workstation, you can: **www.linuxjournal.com/video/run-your-windows-partition-without-rebooting**.

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

OUTRAGED
by the high cost of Fibre Channel or iSCSI Storage ?

Switch to AoE !

SR2461

Coraid Offers a Complete Line of Clustered Modular Storage Products:

- High Performance EtherDrive® SATA+RAID Storage Appliances with 1 GigE or 10 GigE Connections
- Clustered VirtualStorage™ Appliances *(a Revolutionary Logical Volume Manager)*
- Scalable NAS Gateways *(File Sharing with EtherDrive® Storage)*

Coraid's EtherDrive® Storage with AoE is fast, reliable disk storage that's easy to use. And it's much more affordable than iSCSI or Fibre Channel!

- Coraid products use open AoE (ATA-over-Ethernet) block storage protocol, for high performance without the TCP/IP overhead
- With AoE, your shared storage capacity is infinitely scalable – at a fraction of the cost of iSCSI or Fibre Channel storage
- We provide a 3-year warranty and free firmware upgrades on all our products, as well as support from first-rate engineers trained in our technology

Product Excellence Award Winner LinuxWorld Conference & Expo Best Storage Solution

EtherDrive® Storage has a field-proven track record and is 1000+ large data storage customers strong.

CORAID

The Linux Storage People

To learn more about this and Coraid's other products, go online or call

+1.706.548.7200 (toll free: **877.548.7200**)

www.coraid.com

Fast App Launching with GNOME Do

When your panel is littered with application shortcuts or you press Alt-F2 every few minutes to launch a new program, check out a better, faster way to launch your programs—GNOME Do.

KYLE RANKIN

There are many different ways to launch applications on a Linux system. For the longest time, I would alternate between the menu system that came with my window manager and typing the application in a terminal. Honestly, I found that half the time it was faster to launch an application inside a terminal than it was to navigate through a system of menus. At some point in my desktop use, I decided to give the default Ubuntu GNOME desktop a try. GNOME presents at least four main ways to launch applications:

1. Navigate the Applications menu at the top of the screen and find your program.
2. Copy frequently launched applications to the desktop and launch them from the desktop.
3. Copy frequently launched applications to the panel and launch them from the panel.
4. Press Alt-F2 to bring up a command window where you can type the command and press Enter.

I tried each of the four main ways, but I guess I'm a creature of habit, because ultimately, I found myself back to my old ways. When I wanted to launch a program, nine times out of ten, I just would go to an open terminal and type in the command from there. Every now and then, I would navigate the Applications menu. That was my habit, until I discovered GNOME Do. Now I've found I use GNOME Do when I launch the majority of my applications and use a keyboard shortcut, or occasionally, the terminal, for the rest. I don't really even need or use the Applications menu anymore.

Welcome to GNOME Do

GNOME Do is an application launcher tool inspired by the Quicksilver and GNOME Launch Box applications. It is available either as a package in your distribution or you can download the program from the official project page (**do.davebsd.com**). You launch GNOME Do in the background along with your desktop environment, and then press Super-Space to open the GNOME Do window when you want to launch an application (Super is the Windows key on many keyboards). After the window appears, type part of the name for an application; for instance, to open Firefox, type `firefox`. You will notice that the moment you press the F key, GNOME Do chooses an application or other result and refines it as you type. You might need to type only `fi` for Firefox to be displayed (Figure 1). In most cases, there also are alternate choices for your keyword, which you can reveal and select with the up and down arrows (Figure 2).

GNOME Do is a learning program, and as you use it, you will notice that it selects results based on your favorite,

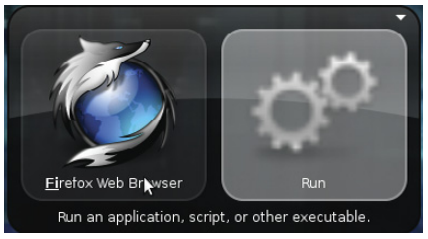

Figure 1. GNOME Do with Firefox Selected

Figure 2. Alternate Keyword Choices

Figure 3. Alternate Actions

most-used choices. This means if you launch F-Spot more often than Firefox, F-Spot shows up first when you press F. GNOME Do also learns which actions (the items that show up in the right pane) you have performed on particular objects and gives those precedence. The ultimate goal is to make it fast and simple to launch applications, open files and interact with different parts of your computer via GNOME Do plugins (more on plugins later in this article).

The interface is easy to grasp, so before too long, you will find it easy to launch programs and interact with plugins and other features.

While the left pane in the GNOME Do window lists objects, the right pane lists actions. An ordinary action for a program like Firefox might be Run, but if you press Tab you will highlight the right pane. Then, you can use the up and down arrow keys to cycle through alternate actions. These actions vary depending on the object, so for an application like Firefox, you might get only the option to copy your typed text to the clipboard or assign an alias. Different objects get a more complete list of actions, so for instance, if the Files and Folders plugin is enabled, it indexes the files in a list of directories. If I start to type a particular filename, it locates matching files. I then can press Tab, and when I press the down arrow key, I will see a number of actions, as shown in Figure 3. In this example, I have the option to open the file, reveal the file in the file manager, move the file to the trash, rename the file or perform a number of other file operations.

Figure 3 also shows that depending on the option you choose, GNOME Do might open a third pane on the right with more options. This often is used when you want to copy or move a file so you can choose its new location.

Preferences

GNOME Do has a pretty shallow learning curve. The interface is easy to grasp, so before too long, you will find it easy to launch programs and interact with plugins and other features. Of course, to get to these features, you first need to go the Preferences window. In the top right-hand corner of the window, you will notice a small triangle. When you click on it, you'll see a drop-down menu with About, Preferences and Quit options. Choose Preferences.

The Preferences window is minimalist (Figure 4) and divides its settings into three tabs. In the first tab, you can change basic settings, such as whether GNOME Do starts at login, whether to show its notification icon and what theme to use. The second tab lets you configure keyboard bindings used with GNOME Do in case you want to change the defaults. The final Plugins tab is

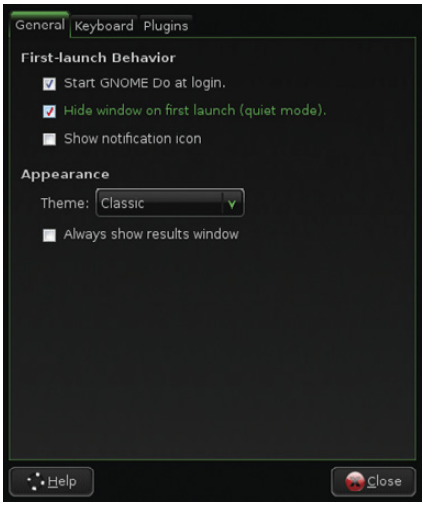

Figure 4. GNOME Do Preferences Window

probably the most interesting. GNOME Do ships with a number of plugins, and additional third-party plugins extend GNOME Do's functionality.

Plugins

GNOME Do's plugins are what move this program beyond a replacement for the Applications menu on your desktop into a blend between a launcher, desktop search tool and central interface for other desktop operations. Most of the plugins probably will be disabled by default, so to get this extra functionality, you need to go into the Preferences window and enable the particular plugin. Below, I highlight a few particularly interesting plugins.

File Plugins

A few different plugins turn GNOME Do into a file browser and search tool. The Files and Folders plugin indexes directories of your choice (highlight the Files and Folders plugin in your Preferences window, and then click Configure to set which folders it indexes). As you type, GNOME Do lists files it finds within those directories as results. You then can copy, delete, browse and perform a number of other options on the files, as shown in Figure 3.

In addition to the Files and Folders plugin is the Locate Files plugin. This plugin uses the GNU locate command, so instead of just searching through directories you specify, you can type a keyword and then select Locate Files in the actions pane. GNOME Do then returns the list of results so you can act on them (Figure 5).

GNOME Tools

With a name like GNOME Do, you probably won't be surprised to know that there are a number of plugins that extend into GNOME functions. The GNOME Dictionary plugin provides a define action, so you can type a word, choose define and get back a definition. The GNOME Screenshot plugin adds a Take Screenshot result if you start to type that phrase. Then you can configure screenshots of the entire desktop, a specific window or even take screenshots after a timed delay. I used that particular plugin quite a bit while writing this article.

The GNOME Terminal plugin extends the traditional Alt-F2 command window in that you can not only run commands within the GNOME Terminal, you also can select particular GNOME Terminal profiles you have created. Figure 6 shows the result when I type `mutt` into GNOME Do. The Open Profile action represents my mutt GNOME Terminal profile.

Finally, the GNOME Session Management plugin gives you the same functionality as the power button at the top of the GNOME panel, so you can lock your screen, shut down, reboot and hibernate your desktop.

Google Tools

A number of plugins can query Google services. The Gmail Contacts plugin indexes your Gmail contact list and provides it as results to queries so you then can select actions, such as e-mail. You also can interface with your Google Calendar using the plugin of the same name and search through and even create new events. The Google Calculator plugin lets you perform the same calculations and conversions you can perform on the Google Calculator site, only within GNOME Do. Finally, with the Google Maps plugin, you can type in an address and select Map to submit the location to Google Maps.

Other Plugins

GNOME Do includes so many great plugins, it's difficult to choose which ones to list. Some other useful plugins allow you to index all of your Firefox bookmarks or Evolution contacts, upload photos to Flickr or ImageShack, update your Twitter status, search through your Rhythmbox music collection, index Tomboy notes and even connect to remote hosts over SSH. And, that's just the list of official plugins. GNOME Do also has many third-party plugins that extend its functionality even further.

I have to say that after trying all sorts of different methods to launch applications, GNOME Do has won me over. If I'm already in a terminal, I still sometimes will launch an application there, but I've switched to GNOME Do to launch most of my programs and even use it as a replacement for a file browser when I want to open a file quickly. If you find that you have a number of key bindings set up to launch programs for you, or if you are tired of reaching for the mouse every time you want to run a program or open a file, I recommend giving GNOME Do a try.■

You also can interface with your Google Calendar using the plugin of the same name and search through and even create new events.

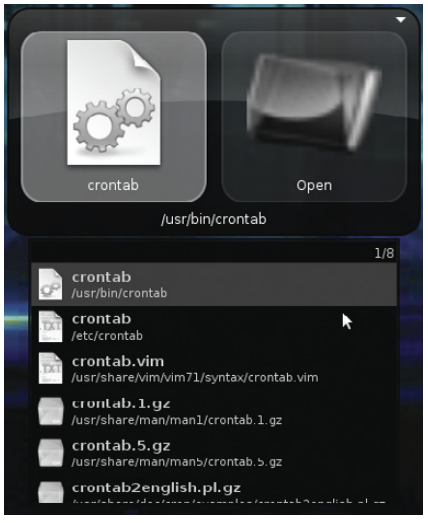

Figure 5. Locate Plugin Results

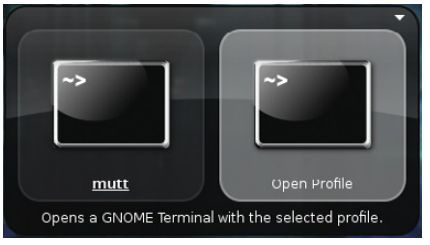

Figure 6. GNOME Terminal Profile Action

Kyle Rankin is a Senior Systems Administrator in the San Francisco Bay Area and the author of a number of books, including *Knoppix Hacks* and *Ubuntu Hacks* for O'Reilly Media. He is currently the president of the North Bay Linux Users' Group.

Software Test & Performance CONFERENCE SPRING

March 31 – April 2, 2009

San Mateo Marriott | San Mateo, CA

Reach Your Peak Performance

The premier technical conference for software testers, QA managers and software developers and managers

Unparalleled Education

- 60+ sessions covering software test/QA and performance issues across the entire application life cycle
- Full-day tutorials offer a complete immersion in the topic

Ultimate Networking

- Share experiences, challenges, ideas, and solutions with your peers and industry leaders — the STPCon faculty

Exciting New Tools & Technologies

- Get a first look at the latest tools and services. Visit with vendors in a relaxed environment. See the products in action, and find the tools and technologies that can help your organization reach its peak performance

Don't miss the best STPCon yet!

Register by March 13th to save up to $300! Use code STP09L

Visit www.stpcon.com to register today!

THE THIRD MAN

If you thought the only desktop options were KDE and GNOME, think again. Xfce provides a full-featured desktop that's lightweight enough to run on old hardware. FEDERICO KEREKI

In the famous 1949 movie *The Third Man*, much hinges on a traffic accident and two men—and, it seems, a third man also was involved. Investigation leads to all kinds of events, until eventually the truth surfaces, and the key third man is found. But, go watch the movie for the details—no spoilers here. In the Linux world, whenever you talk about desktop environments, everyone typically remembers only two (KDE and GNOME), but there's also an often-forgotten third one, Xfce. (There actually are multiple "third" desktops, but let's pretend there's only one for the sake of my catchy intro.)

In this article, I cover Xfce's main features and functions, and why you shouldn't merely dismiss it, because it's a worthy contender to the other more-famous counterparts. Oh, and while you're reading, you might want to listen to the "Third Man Theme", with its distinctive zither sound (not that it has anything to do with Linux, but it's great music).

Xfce started out in 1996, as a Linux version of CDE (Common Desktop Environment), a commercial desktop still in use today. However, after a dozen years of development and several major versions (Xfce currently is at version 4.4, with version 4.6 in the works), Xfce has diverged from CDE and stands on its own. The first versions were based on the proprietary XForms library (see the What's in a Name? sidebar) and were not open source, but version 3.0 was rewritten from scratch, substituting GTK+ for XForms, and was licensed under the GPL. Version 4.0 saw yet another major upgrade, changing to the GTK+ 2 libraries, also used for GNOME.

As its creator Olivier Fourdan said, Xfce is "designed for productivity", so "it loads and executes applications fast, while conserving system resources". With modern hardware, that point may be moot, but Xfce can give new life to older, slower processors or RAM-challenged machines. However, even with the latest CPUs, you might appreciate the extra speed.

All the standard packages included with Xfce (more on this below) were designed with speed and responsiveness in mind, and the rest of the selections also follow suit. For example, instead of other more resource-intensive suites, you get Abiword and Gnumeric—less capable perhaps, but more appropriate given Xfce's goals, and for many users, they're more than sufficient.

Xfce sports no fixed release plan, employing instead the oft-used OSS method of "when it feels like it's ready to be released". The focus is on quality rather than on fixed timelines. Xfce's maintainers also suggest, tongue in cheek, that they can be hired to produce new versions on demand, but it's going to cost you. At the time of this writing, Xfce stands at version 4.4.3, but version 4.6 is in beta and expected to be ready in early 2009.

What's in a Name?

When the project started in 1996, it was named XFce (with an uppercase F), and the letters stood for XForms Common Environment. When XForms' usage was dropped, the name didn't change, but the F became a lowercase f.

Due to its leanness, Xfce also has been called the Cholesterol Free Desktop Environment (but that doesn't quite fit the acronym). The Xfce Wiki site suggests another possibility, X Freakin' Cool Environment, which hasn't caught on yet. So, Xfce is now an acronym that doesn't stand for anything.

Getting Started

You can use Xfce with practically all modern distributions. Some, such as Xubuntu or the Fedora Xfce Spin, come with Xfce as the standard desktop environment. Many others, including OpenSUSE or Slackware, allow you to install Xfce instead of, or in addition to, KDE and GNOME. Generally, you can use a standard package manager to install Xfce. For example, in Ubuntu, you would do `sudo apt-get install xubuntu-desktop`, and in OpenSUSE, you would do `sudo zypper in -t pattern xfce`. As always with open source, you can download, compile and configure Xfce yourself; see Resources for more information.

Additionally, the os-cillation Software Center provides a graphic installation wizard to help with compilation and installation, but it can be a long process. You're better off getting a binary package if you can.

Requirements for Xfce are meager—and well below those of KDE and GNOME—meaning you could run it comfortably on a Pentium III at 133MHz, with 64MB of RAM. Of course, a more powerful processor and a larger amount of RAM will enhance the performance. Some users report running Xfce with even lower-end systems, including a Pentium I or just 32MB of memory, but that's probably the absolute bottom.

If you're running Xfce with Xorg 6.8 or above and an appropriate video card, you can enable several graphic effects. First, make sure you enable the Composite extension, by including the Composite option in the /etc/X11/xorg.conf file, as follows:

```
Section "Extensions"
        Option "Composite" "Enable"
EndSection
```

Then, log out and back in.

You will be able to specify some visual effects, such as transparency for window decorations, window shadows and so on (see below for details). Note that XFWM4, Xfce's own window manager, does the effects on its own, without requiring any further programs or modules.

Customization, Customization, Customization

If you are used to KDE or GNOME, you'll notice some differences, but nothing too dramatic. Right-clicking on the desktop produces a menu with all your applications. Note, however, that the menu is "shallow"—selecting an item in the menu directly produces a list of possible applications, and there are no more submenus. If you want a nicer, multilevel menu, right-click on the menu button on the panel, and select Edit Menu. You'll see a line that looks like this:

```
---include---          system
```

Right-click on it, and you can select the menu style you want: Simple (a single level) or Multilevel. Click File→Save, and then close the window. By clicking on the menu button and selecting Properties, you can manage other changes; feel free to experiment.

If you followed the steps in the above paragraph, you will have experienced Xfce's mouse-only style of configuration. Almost all available options can be selected with only the mouse, providing a consistent and easy interface.

Requirements for Xfce are meager—and well below those of KDE and GNOME—meaning you could run it comfortably on a Pentium III at 133MHz, with 64MB of RAM.

Let's move on to more customizations. In the main menu, choose Settings (Figure 1), allowing you to change Xfce's look and feel. If you search the Net a bit, you'll even find people who have managed to make Xfce look like Windows.

Here are some of the items you might want to explore:

- Appearance (or User Interface Preferences) lets you select or install window and icon themes; see Resources for more eye candy.
- Autostarted Applications lists the applications that will be started automatically whenever you log in. Note that all applications you saved the last time you logged out also will be started; take a look at Sessions and Startup for some options. Also, Preferred Applications lets you specify your favorite Web browser, mail reader and similar programs.
- Desktop Settings and Screensaver let you select the desktop background (solid colors or gradients, or an image) and screensaver, along with some behavior aspects, such as the meaning of a middle- or right-click (the defaults are showing the window list and showing the desktop menu, respectively) or what kinds of icons (if any) will be shown.
- Mixer Settings (or Sound) and Monitor Settings (or Display Settings) deal with sound and screen and have relatively few options. For multimedia options, check Gstreamer

Figure 1. Practically all configuration can be done with only the mouse.

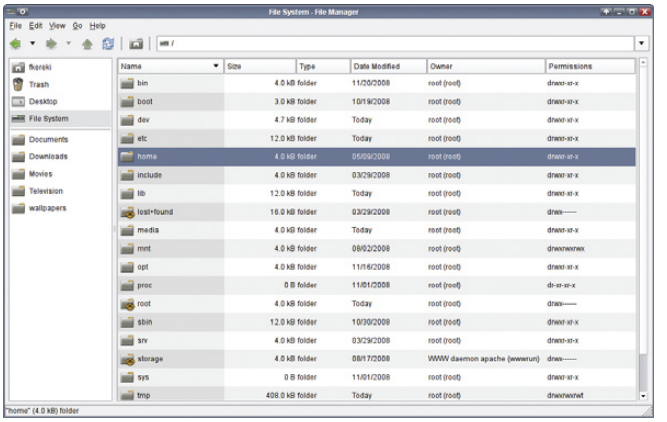

Figure 2. Thunar, Xfce's Standard File Manager

Properties too.

- Keyboard Properties lets you define shortcuts and accessibility features.
- Panel Manager lets you specify how many panels there should be, at which positions (top- or bottom-centered) and of what size. You can right-click on a panel and select Add New Item in order to decide what should be shown. Here, you can change the menu, windows bar and system tray. Opting for a classic look, I configured just one bottom-centered panel with a menu, several program launchers (allowing access to a terminal, editor, file manager and such), a task list (showing open applications), a few applets, a clock and the lock and logout commands.
- Window Manager Settings lets you select the default style for windows, keyboard shortcuts, several details on focusing windows, opaque moves, resizing and the meaning of a double-click on a window. You also should look at Window Manager Tweaks for similar items. In particular, go to the Compositor tab, which lets you specify transparency and shadow parameters.
- Workspace Settings lets you choose how many desktops you want (the fewer the better, in terms of speed) and some other working details.

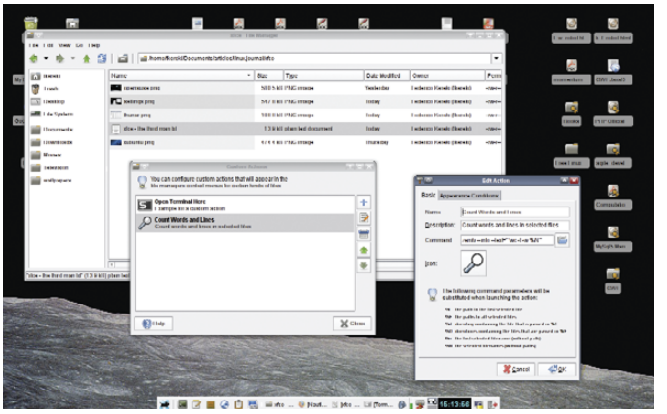

Figure 3. Custom Thunar Action

You might have noticed there is no Fonts configuration option, and apparently, there won't be one in Xfce 4.6 either. If you want to add or remove fonts, you have to do it manually.

Managing Files

Since version 4.4, the default Xfce file manager is Thunar (the old Saxon name for Thor, the Nordic god of thunder), which replaced the previous file manager XFFM (Figure 2).

Thunar is fast and easy to use, and it's similar to Nautilus,

Thunar Customization

Thunar provides extra customization possibilities, allowing you to define personal commands. Click Edit→Configure custom actions, and you can create an action of your own. To test it, I added a "Count Words and Lines" feature (Figure 3). I set Name to Count words and lines, and Command to:

```
zenity --info --text="`wc -l -w %N`
```

And, in the second tab, Appearance, I checked that this command should be applied to Text Files. Zenity shows a dialog box. If you don't use zenity, the output of the wc command won't be shown. If you right-click on a text file and select the Count Words and Lines action, you will get a dialog box showing the result of the wc command.

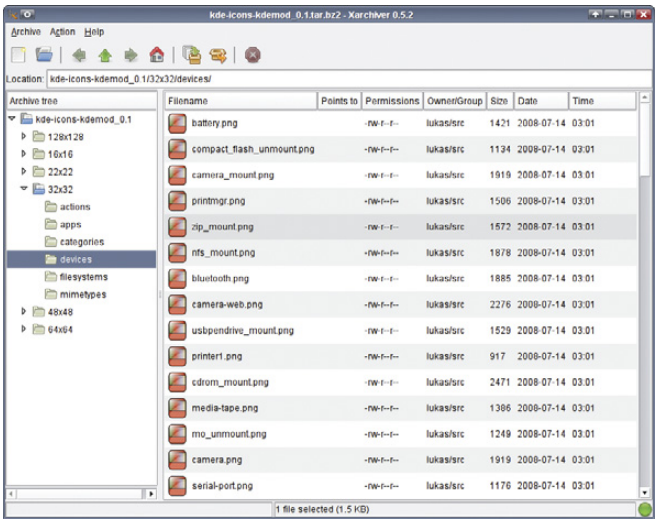

Figure 4. Xarchiver, the Current Xfce Archive Manager

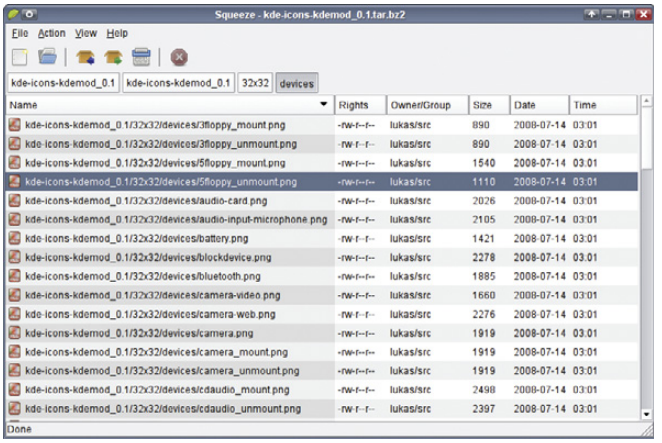

Figure 5. Squeeze, the Xfce Archive Manager

Dolphin or Rox-Filer. In terms of working with files and directories, usage is quite similar to other file managers, and you'll likely feel at home quickly.

Thunar is a lightweight program by design, but you can add functionality through plugins. Download plugins using your package manager or directly from the Thunar Web site. Among the possible extra functions are the following:

- Advanced Properties adds extra pages to the File Properties dialog. For image files, it displays only the image properties, and for .desktop files, it provides launching information, allowing you to specify which program should be run.
- Archive lets you create and extract files from .rar, .zip and similar archive files.
- Renaming lets you rename several files at once and provides search-and-replace patterns, so you could, for example, change all *TXT files to *txt with a single command.
- Media tags also lets you rename media files (such as .mp3) by providing access to their tags.

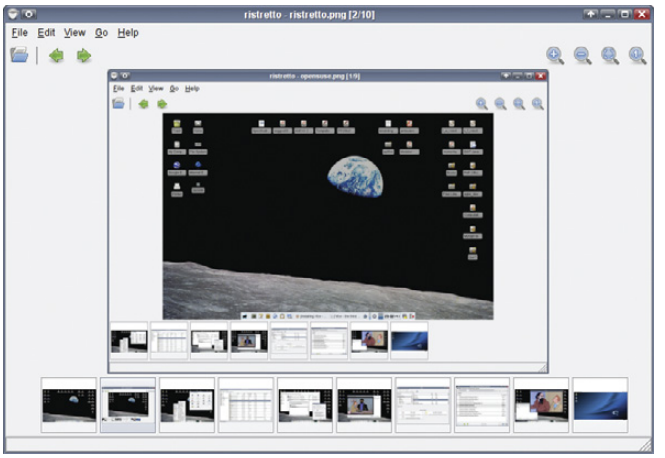

Figure 6. Ristretto Image Viewer

For archived files, the default still is Xarchiver, but Squeeze is set to be the next option (Figures 4 and 5). Xarchiver supports most types of archive files (bzip2, gzip, rar, rpm, tar, zip and so on), with password detection (for reading) and encryption (for writing). You can preview, cut, copy, paste, rename, and drag and drop files to or from archives. Squeeze still is in development and offers only add, extract and delete functions for now.

Finally, Ristretto (the name for a highly concentrated espresso) is an image viewer. You can open a whole directory at once and see all the images in a slideshow fashion (Figure 6). Ristretto lets you zoom and rotate images, and it also can be used via a Thunar plugin.

Instead of more resource-hungry office suites like OpenOffice.org or KOffice, Xfce provides GNOME's Abiword and Gnumeric.

Even More Programs

If you want to play music, Xfce includes Xfmedia (Figure 7). It provides basic playlist functions and is easy to configure. You can randomize the playlist, set diverse "repeat" options as well as provide "visualizations" to accompany the music (not fully developed yet). Xfmedia also is touted as a video player, because of its Xine usage.

However, for DVD playback, it's seriously lacking in command options, such as menu navigation, choosing subtitles and the like, so you're better off sticking with Totem (Figure 8), which is the default application for DVDs.

Xfburn provides CD and DVD creation with a simple interface (Figure 9). Xfburn still is in the early stages of development, but you can burn ISO images or data to CDs or DVDs. The ability to create audio CDs is currently lacking, but it's promised for a future release.

Instead of more resource-hungry office suites like

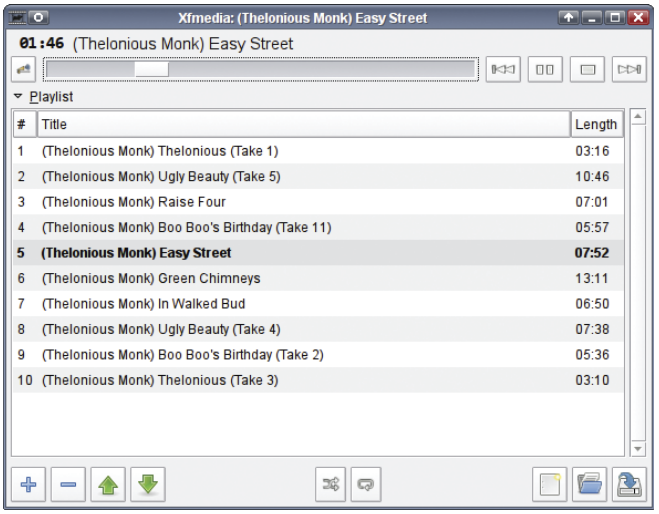

Figure 7. Xfmedia

OpenOffice.org or KOffice, Xfce provides GNOME's Abiword and Gnumeric. Of course, you can use the other suites if you prefer. For simpler text-editing tasks, Mousepad is the standard editor. It offers basic functionality, and in fact, I used it to write this article.

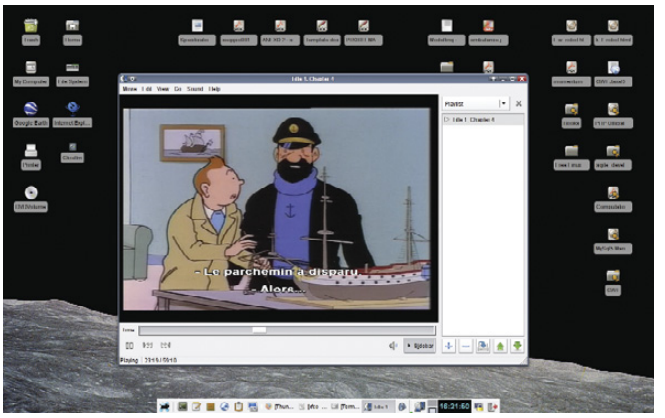

Figure 8. Totem for Video

For messaging purposes, Pidgin also is provided. Pidgin is a good choice, because it can work with multiple protocols (AIM, Google Talk, ICQ, MSN, Yahoo and more), and it integrates well with the desktop.

Finally, as an agenda and calendar, Orage (previously known as Xfcalendar) provides a system-tray clock (right-click on the panel to make it appear) and a personal agenda. You can store events and get alarms. You can

Resources

Xfce Official Web Site: **www.xfce.org**

Xfce Documentation Wiki: **wiki.xfce.org**

Installation Packages: **www.xfce.org/download**

Eye Candy for Xfce: **www.xfce-look.org**

Abiword: **www.abisource.com**

Gnumeric: **projects.gnome.org/gnumeric**

Orage: **www.xfce.org/projects/orage**

Pidgin: **www.pidgin.im**

Squeeze: **squeeze.xfce.org**

Thunar: **thunar.xfce.org/index.html**

Xarchiver: **xarchiver.xfce.org**

Xfburn: **www.xfce.org/projects/xfburn**

Xfmedia: **spuriousinterrupt.org/projects/xfmedia**

CDE: **www.opengroup.org/cde**

Dolphin: **dolphin.kde.org**

GTK+: **www.gtk.org**

Konqueror: **konqueror.kde.org**

Nautilus: **projects.gnome.org/nautilus**

os-cillation: **www.os-cillation.com**

Totem: **projects.gnome.org/totem**

Zenity: **directory.fsf.org/project/zenity**

Distributions with XFCE:

Debian Xfce Group: **pkg-xfce.alioth.debian.org**

Fedora Xfce Spin: **www.redhat.com/archives/fedora-announce-list/2008-February/msg00005.html**

Linux Mint Xfce Community Edition: **www.linuxmint.com/edition.php?id=27**

Mandriva Xfce Live: **wiki.mandriva.com/en/XfceLive**

OpenSUSE Xfce: **en.opensuse.org/Xfce**

Slackware: **slackware.com**

Xubuntu: **www.xubuntu.org**

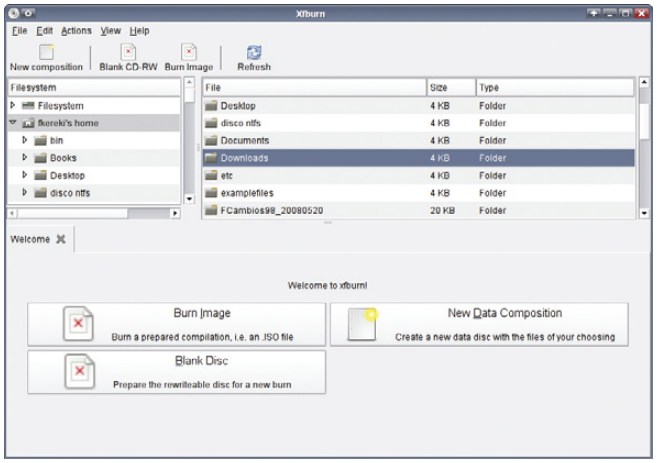

Figure 9. Xfburn provides CD and DVD burning, but it can't create audio CDs yet.

schedule repeating, periodic events and full-day events (like birthdays or holidays). Orage even can work across different time zones (Figure 10).

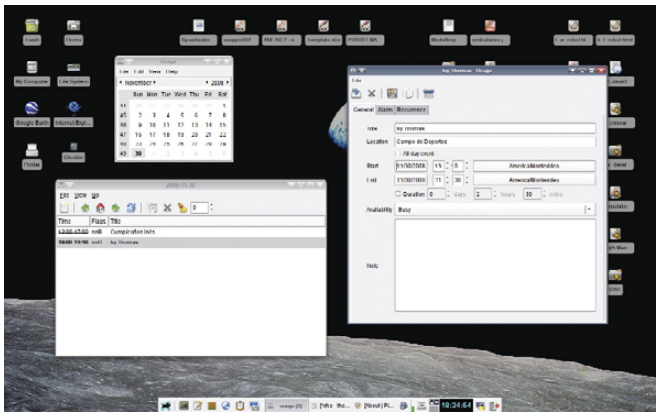

Figure 10. Orage provides clock and calendar functions with alarms and notifications.

Summary

In the movie, finding the third man proved to be a disappointment to the main character, but with Xfce, that won't be the case. With low requirements, high performance and an easy interface, you can be productive in little time. Although it's not quite as packed as KDE and not as simplified as GNOME, it's a balanced desktop environment in terms of power and ease of use. No matter what kind of machine you use, you'll find it worthwhile.■

Federico Kereki is an Uruguayan Systems Engineer, with more than 20 years' experience teaching at universities, doing development and consulting work, and writing articles and course material. He has been using Linux for many years now, having installed it at several different companies. He is particularly interested in the better security and performance of Linux boxes.

***Did you know* Linux Journal *maintains a mailing list where list members discuss all things Linux? Join* LJ's *linux-list today:* http://lists2.linuxjournal.com/mailman/listinfo/linux-list.**

GENSTOR

Linux - FreeBSD - x86 Solaris - MS etc.

Proven technology. Proven reliability.

When you can't afford to take chances with your business data or productivity, rely on a GS-1245 Server powered by the Intel® Xeon® Processors.

Quad Core Woodcrest

2 Nodes & Up to 16 Cores - in 1U

Ideal for high density clustering in standard 1U form factor. Upto 16 Cores for high CPU needs. Easy to configure failover nodes.

Features:

- 1U rack-optimized chassis (1.75in.)
- Up to 2 Quad Core Intel® Xeon® Woodcrest per Node with 1600 MHz system bus
- Up to 16 Woodcrest Cores Per 1U rackspace
- Up to 64GB DDR2.667 & 533 SDRAM Fully Buffered DIMM (FB-DIMM) Per Node
- Dual-port Gigabit Ethernet Per Node
- 2 SATA Removable HDD Per Node
- 1 (x8) PCI_Express Per Node

Servers : : Storage : : Appliances

Genstor Systems, Inc.

780 Montague Express. # 604
San Jose, CA 95131
Www.genstor.com
Email: sales@genstor.com
Phone: 1-877-25 SERVER or 1-408-383-0120

intel Xeon inside™

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

BROWSER BATTLES

OPERA, CROSSOVER CHROMIUM and FLOCK

Firefox is the wunderkind of Linux browsers, but is it right for you? We take a look at three competing browsers—Opera, CrossOver Chromium and Flock—that may just give you more *Webvergnuegen*.

What a relief we felt when Firefox was first released. It replaced "None of the above" as our favorite browser, and we finally were rid of the monolithic dinosaur Netscape. Since then, Firefox has been relatively unchallenged in its supremacy.

As of late, some new challengers are seeking to steal some of Firefox's thunder. In this article, we take a closer look at three of them: Opera, CrossOver Chromium and Flock.

Opera clearly has ambitions, because although always good, recent releases have shown vast improvements, a wealth of smart features and a sleek Euro style. CrossOver Chromium is Google Chrome running on Wine. Though Chromium is betaware, it gives us a sneak peek into Google's plans to re-invent the browser by going minimalist. Will Chrome do to Firefox what Firefox did to Netscape? Finally, Flock seeks efficiency not in raw speed but in doing everything in one place and integrating the Net experience in one "portal". You may find that one of these browsers is a better fit for you.

JAMES GRAY

The Opera Web Browser

Although Linux-friendly for ages, the Opera Web browser, which is based on the proprietary Presto layout engine, has failed to reach critical mass in the collective Linux consciousness. Lacking the massive community involvement and open-source credentials of Firefox, most of us don't even register Opera. In the 2008 *Linux Journal* Readers' Choice Awards, less than 5% of respondents picked Opera as their favorite Web browser (compared to Firefox's 86%).

Although I'd also love to see an open-source Opera, I somewhat understand the company's unwillingness to release its source code. I recently spoke with Opera's CTO, Hakon Wium Lie, who explained the company's position on open source:

> At Opera, we believe in open standards, security, speed, performance and features—these are values that we share with the Linux community....We're very proud of our source code, and we'd like to show it to others, but we haven't found a business model that allows us to do so while still charging for commercial use. Ideally, I'd like to see an open-source license similar to the Creative Commons noncommercial license. The license would say, "here's the source code, feel free to use and reuse it, but we'd like a cut if you make money from it". Anyway, on the Web, I believe open standards are much more important than open source.

Personally, I don't see why an open-source Opera could not only scale up its market share but also leverage that increased popularity to rake in a hefty share of revenues from partnerships with search engines, like Firefox does. However, I'll leave that discussion for another day. The reality is that the Opera browser is good—surprisingly even as good as the Windows edition—which makes it worthy of our scrutiny. Let's have a look.

The Opera Experience

Although I have used Opera on and off over the years as a backup browser, I never really gave it a hard run for its money. When I dove into the new Opera 9.6, I was pleasantly surprised at its features and ergonomically sound look and feel. From a style standpoint, I like Opera better than Firefox on Linux. The Opera folks clearly put much thought into design elements.

Opera makes up for its open-source "deficit" with cool features and customizability. Although the browser is speedy enough, Opera is about the features, not leanness. Many features that are Firefox extensions are already built in to Opera.

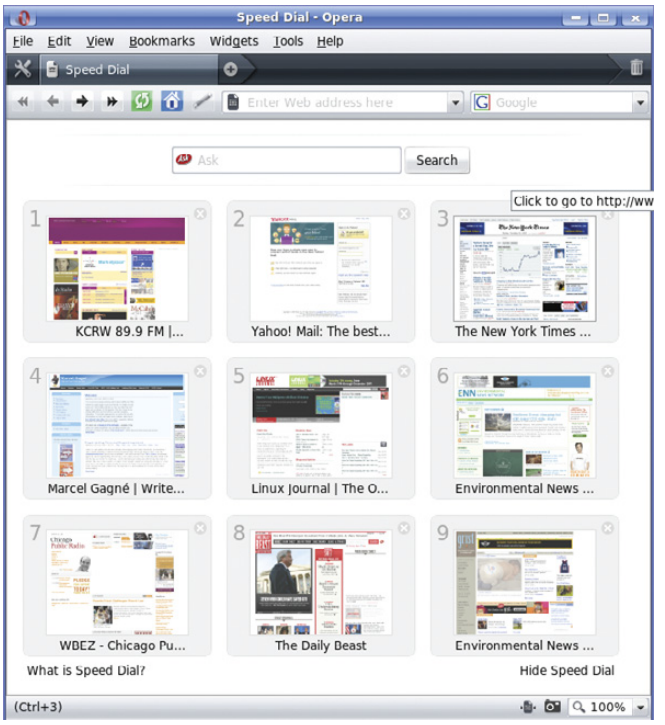

Figure 1. Summon a new tab in Opera, and up pops Speed Dial, a sort of home page with nine one-click links to your favorite sites.

Here are some core features that distinguish Opera from its rivals. The Opera browser has the most interesting startup options of any Linux-based browser. First, upon opening a tab, Opera's default is Speed Dial, a sort of home page with nine customizable thumbnails for your favorite Web sites (Figure 1). Second, Opera lets you decide how to start each session—just as you left off, your home page, a blank page, via dialog or from a saved session. The built-in session manager is a powerful feature, given how many different tabs/windows the typical user has open concurrently. One can have saved sessions for different modes, such as home, work, finances, news, hobbies and so on.

I also found Opera's "philosophy" of default tab-oriented browsing pleasant, which I personally think is nicer than Firefox's tab mode. As you open new pages, Opera opens each substantial (that is, non-squished) tab, complete with mouse-over preview, across the top of the browser. Although you certainly can configure Firefox to act in a similar way, Opera feels more comfortable and looks nicer in tab mode.

A number of other built-in features make Opera worth a look. The Wand allows you to save user names and passwords and autofill them upon subsequent site visits. Opera Link is a service that synchronizes bookmarks, the Speed Dial contents, browsing history and other elements across computers or other devices that run Opera, including mobile devices that utilize Opera Mini, the mobile version of Opera. Furthermore, RSS (with label-able feeds), IRC, a mail client and BitTorrent support are built in.

For those who appreciate aesthetics, Opera goes well beyond its pleasant default skin. Changing Opera's skin is easy, because a surprisingly wide range of skins are easy to preview, download and install with just a few clicks.

Besides the myriad built-in features, Opera also offers a range of widgets that collaborate with the browser. Widgets get their own tab from which one can view, install and manage them. They can appear anywhere the user wants. Although numerous, useful widgets exist, Opera's philosophy is to have more core features built in than does Firefox. The result is a heavier but more immediately customizable browser.

On the negative side, though I didn't find Opera to be noticeably slower than its peers, I had some occasional page-

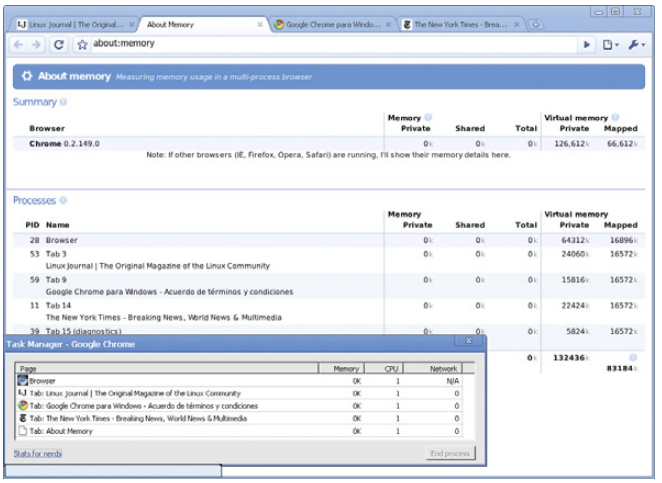

Figure 2. Chromium

rendering problems. On one site, the search box slid behind the main menu and was completely unaccessible.

Despite such problems, I'd argue that Opera's plethora of features makes up for the speed deficits and its rendering problems. Of course, the trade-offs are yours to weigh, as well.

Google Chrome on Wine: CrossOver Chromium

At the time of this writing, the recent Google Chrome browser is available natively only on the Windows platform with a Linux edition still in development. The idea behind Chrome is to remedy the past "mistakes" made by browser makers and provide a tool that also runs applications and not just displays Web pages.

Thanks to good ol' Linux-geek ingenuity, however, we don't have to wait for the Google folks to finish their project. Instead, CodeWeavers, makers of the CrossOver line of Wine-based applications, has created CrossOver Chromium, a Linux "port" of Google Chrome.

The Chromium on Linux Experience

Although the situation may change by the time you read this, the reality is that CrossOver Chromium is betaware and not yet ready for prime time. On Chromium's status, CodeWeavers offers the warning that it is "just a proof of concept, for fun, and to showcase what Wine can do".

Should you decide to tinker with Chromium, you'll find some things satisfactory, and others frustrating. On the plus side, Chromium is thankfully open source, written in V8, Google's JavaScript virtual machine and based on the WebKit layout engine. Furthermore, installing Chromium is a snap, with a downloadable binary in DEB (for Ubuntu) and RPM (for SUSE, Mandriva and Red Hat) formats, as well as a package with a one-click shell script for other distros.

You'll also like Chromium if you prefer a lack of clutter over feature-rich functionality. I've been told that the Windows version is fast as lightning, as reputedly shall be Chrome for Linux, but Chromium certainly is more like thunder than lightning.

One unique feature our ilk will love on Chromium is the (for real) "Stats for nerds" function. When you call forth the task manager, you'll get a list of open Web pages, complete with memory status, CPU usage and network speed related to each page. On the task manager, in addition to an option to kill the process for each open page, you also can click on a Stats for nerds link, which pulls up a new tab complete with additional information such as PID, memory utilization for both the entire browser and each tab.

I also like Chromium's option to open a number of home pages upon startup and not just one. Apparently unique to Chromium is another neat feature, incognito browsing, which allows the user to leave no local trace in the browser or cache of what's been viewed.

On the frustrating side, although you'll find Chromium functional for loading Web pages, it feels a bit slow and clunky. Furthermore, Chromium lacks an integrated non-Web-based RSS reader, which Firefox and Opera have right in the address bar. It also lacks extensive bookmark management, which most other full-featured browsers have. Another core issue is privacy on Chromium, because Google collects usage statistics and crash reports from you as a default. Luckily, you can turn this off by selecting Options from the "wrench" menu.

In the future, Google says we can expect currently absent features to become available, such as better bookmark management, an extension framework, a way to e-mail complete Web pages and links easily and more.

The take-home message on Chromium is that it is a Spartan browser with few bells and whistles whose post-beta experience is slated to be lightning fast on the Linux platform. For now though, CrossOver Chromium is a sluggish prototype held together by lots of virtual duct tape. We can only hope that the native Linux version of Chrome will be so blindingly fast that it gives us a reason to consider leaving its bulkier competitors behind.

Flock

At the far opposite end of the philosophical spectrum from the slim Chromium lies Flock, the Swiss Army knife of browsers. Flock takes the "portal" approach to browsing, adhering to a philosophy that efficiency lies in consolidation rather than raw page-load

Figure 3. The Flock browser seeks efficiency via tight integration with popular sites like YouTube and Facebook.

Figure 4. Flock packs a lot into a small space. The secret to keeping your media organized is the Flock Toolbar. Each icon represents a different type of media by function.

speed. Based on the latest Gecko engine with a number of specialized extensions, Flock bills itself as "The Social Web Browser" that "collects all of your feeds, friends, media and sites in one convenient place".

The reason for such a billing is that Flock wants you to do nearly everything from a unified command center. This includes interacting with social networking sites (such as Facebook and Twitter), photo and video search (such as Flickr and YouTube), photo uploading and sharing (such as Picasa and Photobucket), blogging (such as Blogger and LiveJournal), news consolidation with custom RSS feeds and bookmark syncing with on-line bookmark services (such as delicious and magnolia).

The key to keeping all this stuff straight is the Flock Toolbar, which offers an icon that represents each type of media by function (Figure 4). These icons include My World, a home base of sorts that includes the information you want, such as news feeds and your Facebook friends; the People Sidebar for social-networking interactions; the Media Bar for quick video and photo searches; the Feeds Sidebar for managing RSS; Webmail for interfacing with Web-based e-mail services; Favorites (bookmarks); the Accounts and Services Sidebar for managing accounts and logins; the Web Clipboard, a location to which you can drag links, images and text to save for later perusal; the Blog Editor and the Photo Uploader. Besides all this, you can, of course, simply surf the Web conventionally.

Happy Flocking

I first approached Flock 2.0 with my own Firefox-colored assumptions and habits, including a primal urge to summon each site I visit onto its own tab or window. I typically scroll through my tabs with Ctrl-Tab or my windows with Alt-Tab. As I dug in to Flock, I had to tell myself "Hold off on the keystrokes and start flocking", as the animated introduction suggests.

As I began to "flock" on Flock, I realized the tight integration with its partner sites. To test Flock's capabilities, I sent a friend a hilarious Bollywood music video I had seen recently on YouTube. Culling my urge to call up YouTube on its own tab, I instead did it the Flock way by launching the Media Bar, which popped up one-inch wide across the top of the browser window. Because the Media Bar is integrated with several media-based sites, including YouTube, I could choose YouTube from the drop-down menu and search the site without going there directly. The Media Bar came back with thumbnails and mouse-over previews of the search results. Then—here's the coolest part—after finding my video, I sent it to my friend by dragging the thumbnail over to his Facebook entry in the People Sidebar, which automatically composed a message to him, including the link to the video and the thumbnail. All I had to do was click Send. I could have done the same with my Yahoo Webmail, Twitter or blog entry. Similar drag-and-drop functionality and integration works while blogging on one of the supported blog sites.

Despite Flock's toolbar-driven modular layout, you're probably wondering how it packs so many goodies into such a small space. Admittedly, all that content was a bit scrunched on my 12"-laptop display, but it still was functional. The real estate found in a large LCD is more appropriate for flocking. Nevertheless, features such as the rapidly sliding tab bar allow you to open and manage a huge volume of tabs.

Otherwise, Flock was very customizable according to my whims, and the main toolbar was logical and functional after becoming accustomed to it. The degree of integration with other sites is unprecedented among Linux-based browsers. Despite the heft of features, Flock is based on Firefox, which means the options are familiar, browsing is nimble, and most, but

not all, extensions are usable.

Whether you choose Flock really depends on the degree to which you hang out on social media sites, share pictures and video with friends, blog frequently and track RSS feeds. If you are a social media addict, I suspect Flock will save you time and hassle. If you are a more casual user, Flock may seem claustrophobic and cluttered. I, for instance, am a Facebook user who checks the site once or twice a day—probably not enough to warrant using Flock. However, I admit that having my Facebook (and Twitter) contacts right next door makes me more conscious of and interactive with my contacts. Who knows, maybe I'll stay on board. If you do decide to flock, however, be prepared to throw your typical browser habits out the window and re-orient yourself to Flock's all-in-one philosophy.

Before putting Flock to bed, I should add that Flock makes two special editions for the Linux platform: the Gloss and Eco editions. The Gloss edition is preconfigured for entertainment and fashion-related topics, and the Eco edition is for green topics. Being an eco-geek, I checked out the latter. On the negative side, the souped-up editions are a release behind—namely 1.2.6 at the time of writing, compared to 2.0 for the standard release. The Eco edition, beyond its earthy eye candy, preloads a plethora of enviro-oriented links, media streams, RSS feeds and favorites. As an avid reader of green media, I was impressed with the wide range of selections, many of which I had never seen before. Although the Eco edition is probably overkill on quantity, it provided me with plenty of new information sources, as well as a template for how to maximize Flock.

Other Browsers for Linux: Epiphany and Konqueror

The Linux platform is blessed with several other Web browsers. Here is a quick take on two, Epiphany for GNOME and Konqueror for KDE.

If you are an avid Ubuntu user, you probably are familiar with Epiphany, the GTK-based Web browser built for the GNOME desktop. Besides integrating tightly with the GNOME desktop, Epiphany's goal is to be simple and easy to use. The browser utilizes Mozilla's Gecko layout engine and offers nearly the same functionality as Firefox, including its extensions. However, one feature that stands out in Epiphany is its topic-based, rather than hierarchical, bookmark management, which is similar to Gmail's labels. This system allows you to categorize a bookmark more intuitively with multiple topics. Epiphany also supports cookie management, pop-up blocking, tabbed browsing and its own extension package. Some native extensions relate to mouse gestures, a certificate viewer, an interactive Python console and smart bookmarks.

If you're a KDE aficionado, you know Konqueror well as your "everything tool". Beyond managing and viewing files, Konqueror also is a decent and basic Web browser, though not as robust functionally or stylistically as Firefox or Opera. I find Konqueror to be a great backup Web browser when things go awry with others. The browser identification tool aids the troubleshooting process by letting you configure how Konqueror reports itself—for example, as Internet Explorer, Googlebot, Firefox and Safari, among others.

Let Your Philosophy Decide

With the above and more options available, what are you going to choose as your browser? We are fortunate that each browser niche is being filled with great options on the Linux platform. For those who believe that efficiency lies in integration and feature-richness, Flock and Opera are excellent options. Flock integrates the most tightly with other services, and Opera builds in as much functionality as possible. Neither browser is known for its page-load speeds. Meanwhile, browsers like Firefox and its more streamlined cousins (such as IceCat and Epiphany) hold the middle ground, seeking to balance speed with essential functionality, leaving many features to the extensible extension system. Currently, Firefox is the fastest we've got. However, although CrossOver Chromium is not so nimble, it gives us an enticing preview of the forthcoming Google Chrome for Linux. Chrome for Windows is faster than Firefox. If you pride speed and agility above all else, keep your eye out for Chrome. Whatever your philosophy, you'll find a great Linux-based browser to meet your needs.■

James Gray is *Linux Journal* Products Editor and a graduate student in environmental sciences and management at Michigan State University. A Linux enthusiast since the mid-1990s, he currently resides in Lansing, Michigan, with his wife and cats.

TECH TIP Running Complex Commands with sudo

If you use sudo to run commands as root, you've probably run into "permission denied" problems when only part of a pipeline or part of a command is running with root permissions.

This fails with "permission denied" because the file is writable only by root:

```
$ echo 12000 > /proc/sys/vm/dirty_writeback_centisecs
```

But, this fails too:

```
$ sudo echo 12000 > /proc/sys/vm/dirty_writeback_centisecs
```

Why? The /bin/echo program is running as root, because of sudo, but the shell that's redirecting echo's output to the root-only file is still running as you. Your current shell does the redirection before sudo starts.

The solution is to run the whole pipeline under sudo. There are a couple ways to do it, but I prefer:

```
echo "echo 12000 > /proc/sys/vm/dirty_writeback_centisecs" | sudo sh
```

That way, I can type everything before the pipe character, and see what I'm about to run as root, then press the up arrow and add the `| sudo sh` to do it for real. This is not a big deal for short, obvious pipelines, but when you're building up a more complicated command as root, it's safer to look at it first before you run it.

—DON MARTI

DrupalCon is the event for Drupal Developers.

It's where decisions are made, clients and talent are found, and new stuff is unveiled.

Drupal is the popular open source content management system that powers some of the most important websites and applications in the world.

Wednesday	Thursday	Friday	Saturday
Showcase + Strategy	Showcase + Strategy	Showcase + Strategy	Code Sprint
Design + Usability	Drupal Comunity	Business Side	Documentation Sprint
Performance + Scalability	Code + Development 1	Code + Development 1	Drupal Dojo
Code + Development	Code + Development 2	Code + Development 2	

★ ★ ★ ★ ★

REGISTER TODAY

www.drupalcon.org

★ ★ ★ ★ ★

AIR
Out Your Desktop with Adobe

Cross-platform never looked so Flashy.

SHAWN POWERS

Most Linux users are intimately familiar with programs that are similar to their commercial counterparts, but not quite the same. There isn't anything wrong with clone applications, but if programs were inherently usable, regardless of the underlying operating system, fewer programmers would have to re-invent the wheel over and over again.

Cross-Platform Apps, a Great Concept

One real hope for Linux users is the idea of cross-platform applications. It's not a new concept, but we're beginning to see some huge leaps in the ability for Linux users to use the same applications that Windows and OS X users use—not just clones, but the same actual apps.

Java, for instance, is a great way to make the underlying operating system unimportant. One really huge stumbling block with Java programs, however, is that packaging them for the different computer systems is very difficult. Sure, you can download a JAR file, but that does the average user little good. It takes effort to make installers for each platform you want to support. Java also has a reputation for poor performance. Before any Java developers form a lynch mob and hunt me down with pitchforks, note that I said it has a *reputation* for poor performance. At one time it was true, but in most cases, these days Java programs perform quite well. However, fair or not, the general opinion regarding Java programs to date is that they are slow and cumbersome.

On the Web, Computers Become Irrelevant

Web-based applications recently have become the most prevalent way to provide equal access for everyone. Whether you call it Web 2.0, advanced JavaScripting or just the availability of a more diverse set of Web programming frameworks, the Web currently is a hotbed of new platform-agnostic programs. That trend isn't going away any time soon, but there are a few problems that are tough to solve with on-line apps. For one, relying on the Web browser to handle multiple applications puts all of your eggs in one basket. If the Web browser itself crashes, so does every one of your running programs. An even

Equal-Opportunity Adobe?

Although the initial development of Adobe's AIR for the Linux platform has been behind its Windows and Macintosh counterparts, Adobe promises future releases all will come out at the same time for all three platforms. I contacted Rob Christensen, Adobe AIR Senior Product Manager, and he confirmed that future releases are planned to come out simultaneously.

At the time of this writing, that means while AIR 1.5 is available for Windows and OS X, the latest version for Linux is 1.1 Beta. The unfortunate side effect of the different versioning is that many of the newest (read: coolest) AIR applications don't run under Linux, because they require the 1.5 runtime environment. In fact, due to the "beta" aspect of the Linux port, some apps designed to work with version 1.1 don't even function properly.

Hopefully, Adobe will remain true to its promise, so AIR apps work everywhere, all the time, regardless of the underlying OS. I'm hopeful, especially after seeing the recent timely releases of Adobe Flash.

Advertiser Index

Advertiser	Page #	Advertiser	Page #
1&1 Internet, Inc. www.oneandone.com	15, 17	Logic Supply, Inc. www.logicsupply.com	41
Aberdeen, LLC www.aberdeeninc.com	37	Microway, Inc. www.microway.com	C4, 29
ASA Computers, Inc. www.asacomputers.com	79	Mikro Tik www.routerboard.com	7
Cari.net www.cari.net	61	O'Reilly MySQL Users Conference www.mysqlconf.com	71
Community One developers.sun.com/events/communityone	35	Polywell Computers, Inc. www.polywell.com	5
Coraid, Inc. www.coraid.com	47	Rackspace Managed Hosting www.rackspace.com	C3
DrupalCon DC dc2009.drupalcon.org	63	Servers Direct www.serversdirect.com	3
Ecomm Media, Inc. ecommconf.com	69	SharePoint www.sptechcon.com	51
Emac, Inc. www.emacinc.com	21	Silicon Mechanics www.siliconmechanics.com	23, 27
EmperorLinux www.emperorlinux.com	9	Technologic Systems www.embeddedx86.com	11
Florida Linux Show www.floridalinuxshow.com	75	Techweb Live Events (BlackHat) www.blackhat.com	39
Genstor Systems, Inc. www.genstor.com	57	Ubiquiti Networks, Inc. www.ubnt.com	C2
IT360 www.it360.ca	25	ZT Group International www.ztgroup.com	1

ATTENTION ADVERTISERS

June 2009 Issue #182 Deadlines
Space Close: March 23; **Material Close:** March 31

Theme: Readers' Choice Awards

BONUS DISTRIBUTIONS:
Sun's Java One, O'Reilly Web 2.0 Summit, ISPCON Spring RailsConf, phptek

Call **Joseph Krack** to reserve your space
+1-713-344-1956 ext. 118, e-mail joseph@linuxjournal.com

bigger problem is that if the Internet itself isn't accessible, neither is the Web-based program.

Several potential solutions exist to solve the "unconnected" problem regarding Web apps. Google, for instance, continues to develop its Gears infrastructure that allows people to use Web applications when in off-line mode. The concept, whether implemented by Google or someone else, will make Web applications more viable as desktop program alternatives. Unfortunately, it still tethers us to a Web browser. Projects like Prism can allow for separate instances of Web applications, but regardless of what browser is used, the apps still require the browser in order to function.

Enter: AIR

Adobe has taken its long history of Flash on the Web and given developers the ability to create standalone applications that don't depend on a Web browser at all. I know that many *Linux Journal* readers just rolled their eyes at the mention of Flash technology under Linux, but to be fair to Adobe, it has put more serious effort in its Linux ports recently than ever before. The mere fact that the same version of Flash is available for Linux as is available for Windows proves that Adobe is taking our favorite operating system more seriously.

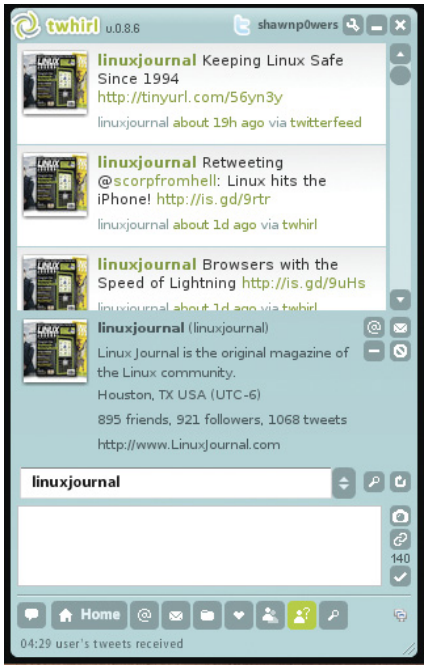

Figure 1. Twhirl is a Twitter microblogging client.

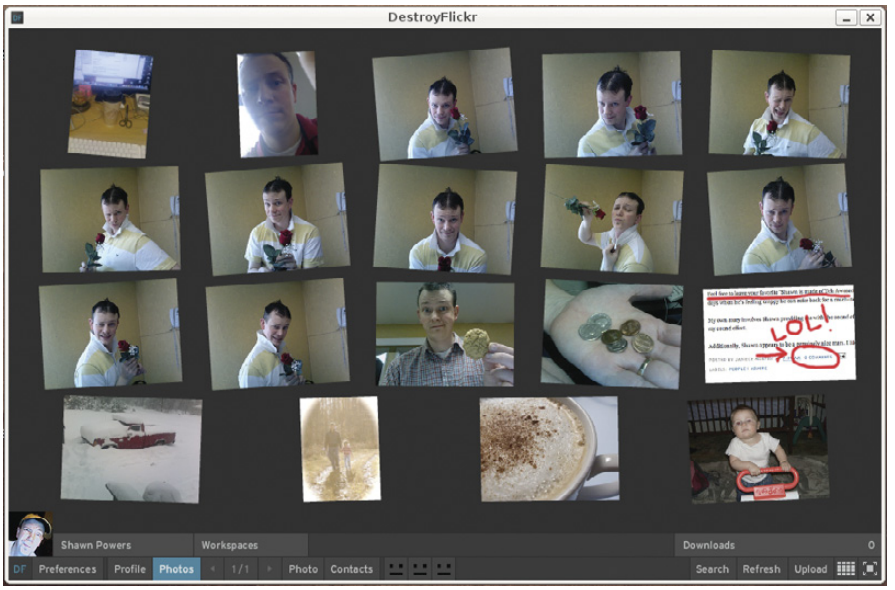

Figure 2. DestroyFlickr manages your on-line Flickr account.

So, what makes AIR unique? Several things:

- Applications look and function the same, regardless of the operating system.
- Developers do not need to package AIR apps separately. One package installs identically on any platform.
- AIR applications, along with the AIR environment itself, can be installed directly from a Web link inside a Web browser (assuming a recent version of Flash is installed on the computer).
- Applications are standalone and don't require a browser. One AIR app doesn't affect other apps if it crashes.

One of the best things about AIR applications is that they tend to look aesthetically pleasing. Let's look at a few from Adobe's AIR Marketplace. I specifically focus on those that work under Linux now, and a couple that I hope work by the time you read this article.

Twhirl

Twhirl is one of the dozens of Twitter clients available. Many people find Twitter's Web interface much less useful than using a dedicated client. I'm in that boat. Twhirl has lots of seemingly simple features that make it a great way to interface with the Twitter universe.

Many users prefer another AIR-based Twitter app, known as TweetDeck. As both are free, and both work well under Linux, so I suppose it's only fair to mention both. Twhirl is just my personal preference.

Twhirl—works under Linux AIR 1.1 Beta: yes.

DestroyFlickr

DestroyFlickr is a program that lets you manage your Flickr stream with an interface that resembles a light table. It's a convenient way to use Flickr.

Another application that currently works under Linux is Flump. It is much more simplistic in its interface, but it can upload and download photos.

DestroyFlickr—works under Linux AIR 1.1 Beta: yes.

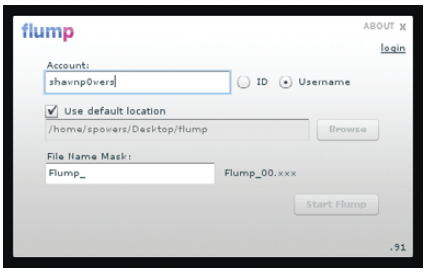

Figure 3. Flump is a very simple Flickr app.

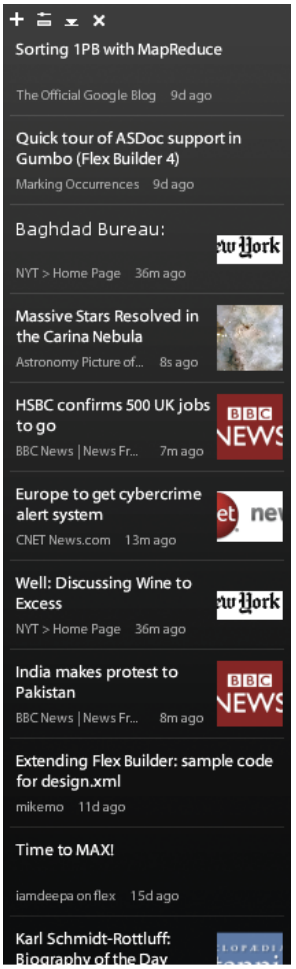

Figure 4. Snackr continuously scrolls your RSS feeds, so you can "snack" on them at your leisure.

Snackr

Snackr is an RSS ticker that runs along the border of your screen. Clicking on a story opens it in a preview window, which gives you the option to open the actual page in a browser.

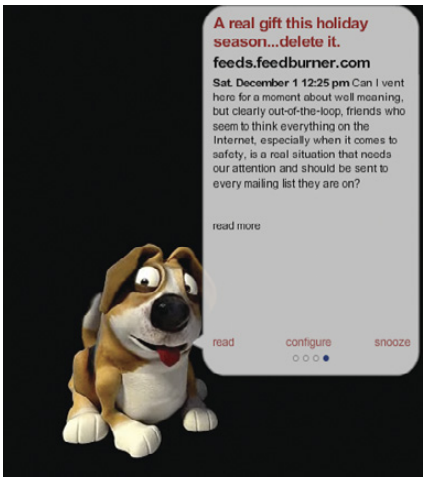

Figure 5. Pet-It RSS uses an avatar much less annoying than Microsoft's Clippy.

If you would prefer to have a cute fuzzy puppy read your RSS feeds to you while he rolls around on your desktop, you might prefer Pet-It RSS News by zerofractal.

Snackr—works under Linux AIR 1.1 Beta: yes.

Snippage

Snippage allows you to select a specific section of a Web site and make it display in a widget on your desktop. The Web site can update automatically, so if the page changes, so does the widget.

Figure 6. Snippage lets you create a widget from a Web page selection.

Snippage—works under Linux AIR 1.1 Beta: yes.

Google Analytics Reporting Suite

If you look at Google Analytics information very often, this application will save you time. Its functionality isn't much greater than visiting Google's Web site, but the speed and convenience is nice.

Google Analytics Reporting Suite—works under Linux AIR 1.1 Beta: yes.

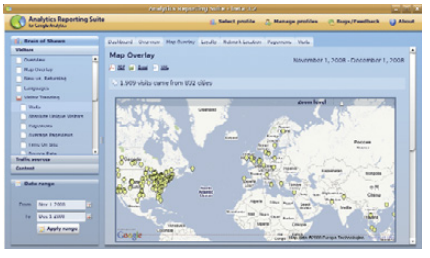

Figure 7. Browsing through Google Analytics information is simple with this application.

Back It!

Back It! is a simple utility for backing up a predetermined set of files and folders from one place to another. The program isn't complex, but it's a great example of AIR interacting with the underlying operating system.

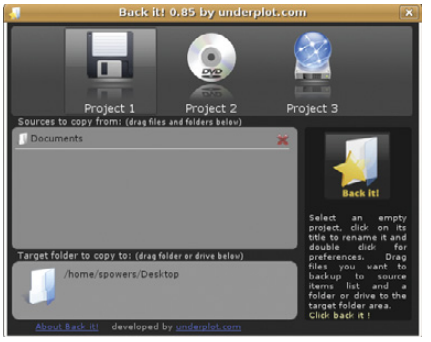

Figure 8. Back It! does exactly as its name suggests. It backs stuff up.

Back It!—works under Linux AIR 1.1 Beta: yes.

Remember the Milk Notifier

RTM Notifier logs in to your on-line account and notifies you of upcoming tasks. I noticed a few graphic glitches running it under Linux, but the program itself is fully functional.

Figure 9. Remember the Milk Notifier is a popular on-line task manager.

RTM Notifier—works under Linux AIR 1.1 Beta: mostly.

Sam's Interactive Reader

This is a children's program that reads with young readers interactively. There also are activities and the ability to download additional content. Pricing for additional content varies from free to slightly more expensive than free.

The functionality of this program is not consistent. Sometimes it won't load

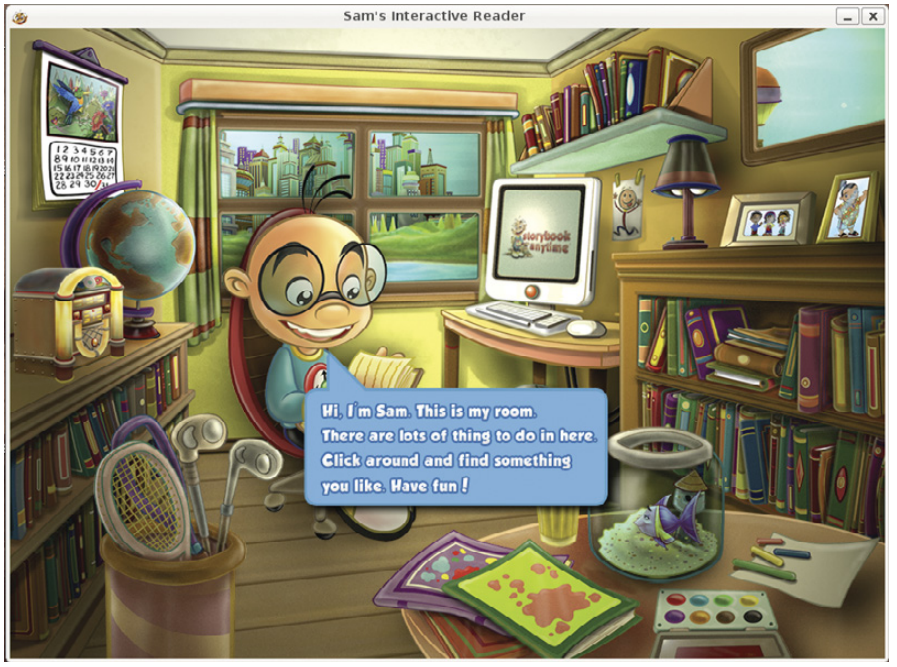

Figure 10. I'm a sucker for childhood literacy, so I really hope this program works better by the time you read this.

at all, and sometimes it starts almost enough to use. I included the app, hoping it works well when AIR for Linux is no longer in beta.

Sam's Interactive Reader—works under Linux AIR 1.1 Beta: almost, but no.

MyMediaPlayer

MyMediaPlayer is an application that interfaces the hulu.com Web site and makes it easy to navigate and display hulu videos.

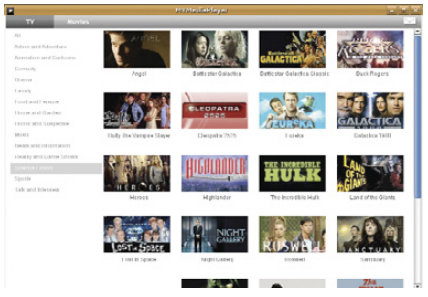

Figure 11. MyMediaPlayer makes the already-simple hulu.com even simpler—assuming it works by the time you read this.

Under the Linux beta of AIR, this application does everything but actually play the videos. The menu navigation is easier than using the hulu.com Web site, and I expect it to be a great application once AIR gets out of beta.

MyMediaPlayer—works under Linux AIR 1.1 Beta: almost, but no.

Pandora Desktop

If you use Pandora to listen to music, you know how inconvenient it can be if you accidentally close the browser window. Several Pandora AIR applications exist in the wild, and assuming the actual playback works once AIR is out of beta, this application will be awesome. It even integrates in the Linux notification area on the taskbar.

Figure 12. Pandora is a Web application just begging to be a standalone app.

As with the other apps that play back media, Pandora apps don't quite work yet. By the time you read this, it should be an application you won't want to forget.

Pandora Desktop—works under Linux AIR 1.1 Beta: almost, but no.

FotoBooth

FotoBooth is an application written in Flex that allows you to take photos with your Webcam. It supplies real-time effects you can apply to the photos and allows for uploading directly to Flickr.

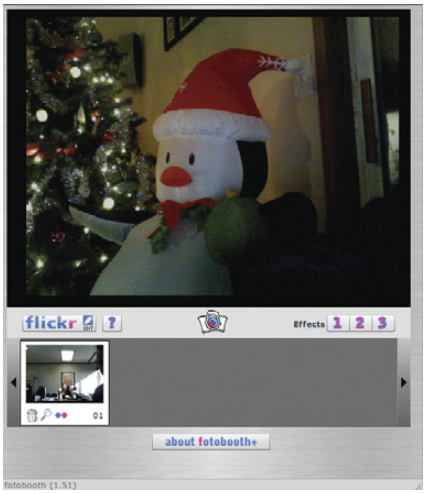

Figure 13. FotoBooth is a clone of Apple's PhotoBooth, with Flickr integration.

FotoBooth has a complex history. It exists as a Web-only Flash application and several versions written in AIR. Some of the versions work under Linux, and some don't. Again, hopefully by the time you read this, it will be a moot point.

FotoBooth—works under Linux AIR 1.1 Beta: yes.

Almost Equal-Opportunity Development

Because countless numbers of AIR apps are available, I'm certain I missed many that you would appreciate. The best place to search for new AIR applications is the Adobe AIR Marketplace (see Resources).

There is a significant chance that you'll want to create your own custom program as well. Fear not, because Adobe also offers tools that allow developers to create their very own AIR apps. Granted, the tools available for Linux developers aren't as robust as those available for Windows and Macintosh users, but the upside is that the Linux tools are free—well, at least for now.

Folks familiar with the Eclipse IDE can download the free (currently alpha, currently free, although that might change) Flexbuilder plugin. It interfaces with the extremely well-known Eclipse program to give Linux users a method

eComm™

EMERGING COMMUNICATIONS CONFERENCE

DEFINING THE POST-TELECOM ERA.

SAN FRANCISCO AIRPORT MARRIOTT · **MARCH 3-5, 2009**

We are **changing** the way the world communicates by **building community**, **providing inspiration** and **spreading the vision** of the post-telecom era.

http://eCommConf.com

REGISTER NOW

LEARN TO HACK:
Telecom Networks & Mobile Handsets

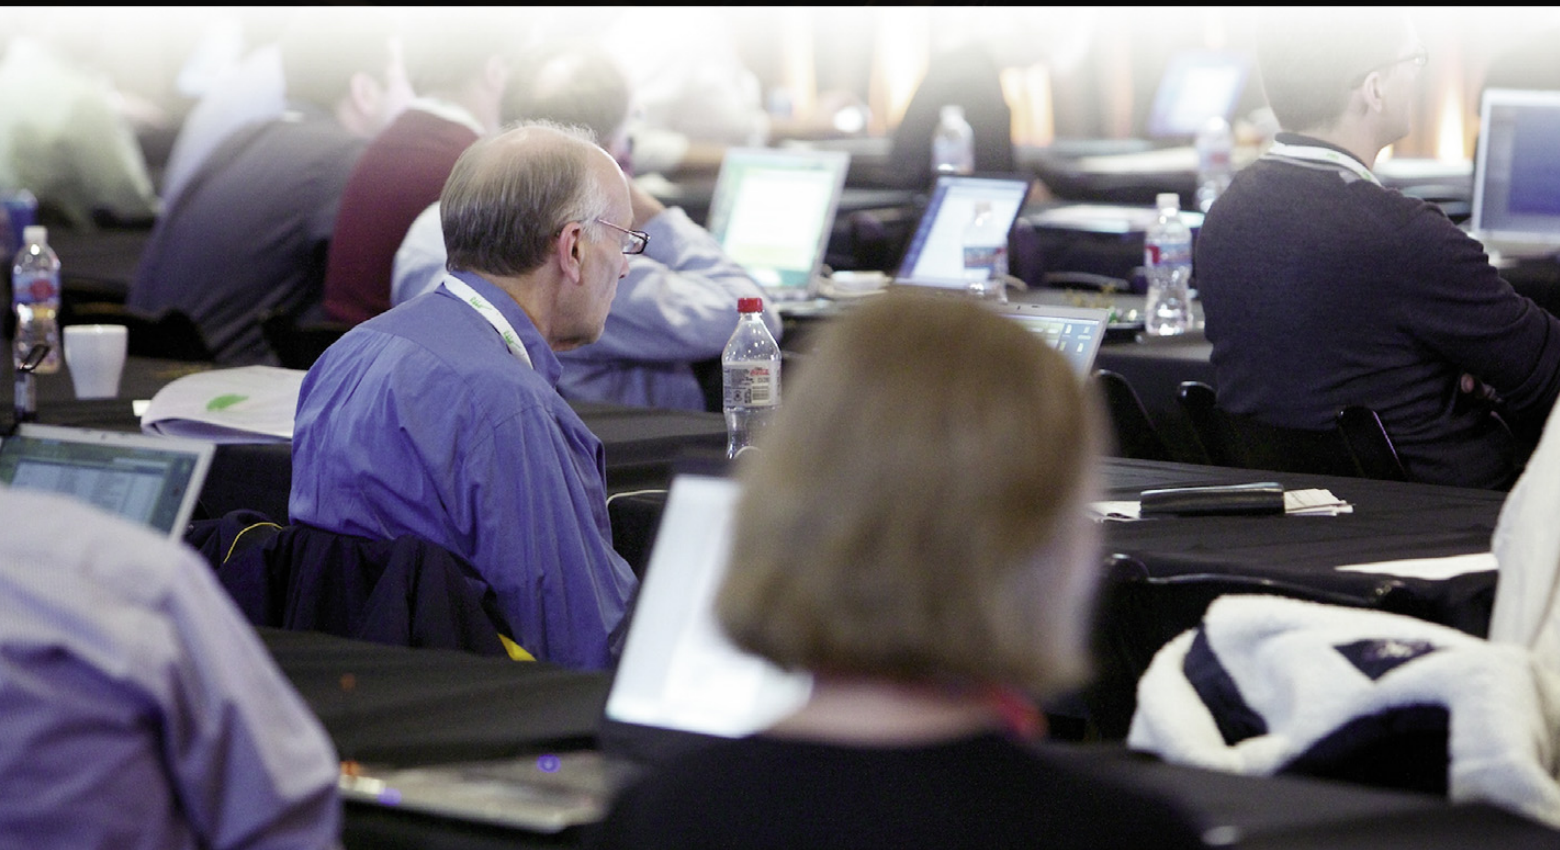

for creating AIR apps. If GUI programming environments don't impress you, AIR apps also can be created with a simple text editor.

When you add the simplistic installation method, the cross-platform homogeny and the availability of (albeit rudimentary) development tools for Linux, Adobe AIR is a platform that levels the playing field a bit for those who prefer Linux. Most encouraging of all, at least for me, is that Adobe is paying attention to Linux in a way that has never before been witnessed. And, that kind of Flash really gets my attention.■

Shawn Powers is on the editorial staff for *Linux Journal*, a technology director for a school district, a dad, a husband and is beginning to lose his hair. He's very happy about all but one of those facts. Reach him via e-mail at shawn@linuxjournal.com.

***LJ* pays $100 for tech tips we publish. Send your tip and contact information to techtips@linuxjournal.com.**

Resources

Adobe AIR Linux FAQ: **tinyurl.com/airfaq**

Adobe AIR Marketplace: **tinyurl.com/airmarket**

Adobe Flexbuilder Linux Page: **labs.adobe.com/technologies/flex/flexbuilder_linux**

Twhirl: **www.twhirl.org**

TweetDeck: **www.tweetdeck.com**

DestroyFlickr: **www.destroytoday.com**

Snackr: **www.snackr.net**

Snippage: **snippage.gabocorp.com**

Google Analytics Reporting Suite: Available on the Adobe AIR Marketplace (see above)

Back It!: **backit.underplot.com**

Remember the Milk Notifier: **rtm-notifier.com**

Sam's Interactive Reader: **www.storybookanytime.com**

MyMediaPlayer: **www.paulyanez.com/labs/mymediaplayer**

Pandora Desktop: **www.pandora.com/desktop**

FotoBooth: **tinyurl.com/fotobooth**

Subscribe Now for JUST $29.50

The 5th annual **MySQL Conference & Expo,** co-presented by Sun Microsystems, MySQL and O'Reilly Media, happening April 20 – 23, 2009 in Santa Clara, CA, brings over 2,000 open source and database enthusiasts together to harness the power of MySQL and celebrate the huge MySQL ecosystem.

Take advantage of this rare opportunity to meet face to face with a huge core group of engineers who are in the process of developing MySQL. Share ideas and get your toughest questions answered from MySQL experts. Bring your team, your laptops, and leave with new insights in how to build high performance scale-out applications using MySQL.

The quality of the sessions and tutorials at the **MySQL Conference & Expo,** is very high. Year after year, the most frequent comment the conference team receives from attendees is that it is sometimes very difficult to choose which great simultaneous session to attend. From novice to expert-level sessions, developer to DBA, business to technical, the sessions provide expert insight into solving architectural, design, and development issues your organization faces.

MySQL® Conference & Expo

APRIL 20–23, 2009
SANTA CLARA, CA

Register Now and Save 15% when you use discount code: mys09lj
www.mysqlconf.com

©2009 O'Reilly Media, Inc. O'Reilly logo is a registered trademark of O'Reilly Media, Inc. All other trademarks are the property of their respective owners. 80608

The Day the Earth Stood Still

Linux visual effects artists around the world create a new sci-fi classic.

ROBIN ROWE

The Day the Earth Stood Still is a re-invention of the 1951 science-fiction film classic. Keanu Reeves stars as the benevolent visiting alien Klaatu, come to Earth to warn us to change our barbaric ways or face destruction.

Ten years ago, *Titanic* was the first film to use Linux in a big way. Today, Linux dominates big-budget visual effects and 3-D animation. Ever since *The Matrix*, it's become routine to have several visual-effects companies working on the same film. A visual effects supervisor at the studio, in this case Fox, selects which companies will create the visual effects.

Twentieth Century Fox

"I came in and met with the director Scott Derrickson", says *The Day the Earth Stood Still* Visual Effects Supervisor Jeffrey A. Okun. "In Scott's opinion, and one I agree with, the day of visual effect as star of the movie is gone. He wanted to focus on story. He wanted spectacular effects that were invisible. When dealing with spaceships, aliens and giant robots, that's a bit of a challenge."

"Weta was our primary group on the film that did 220 shots on the film", says Okun. "Then Cinesite. We had Flash Filmworks and CosFX. Later on we added Hammerhead and Hydraulx, a company called At the Post, and a couple other little companies. Weta handled the Sphere, the alien, the robot and the Swarm. It's all particle systems based on chaos theory. That means it's render-intensive."

"There's a shot of the Sphere that we call the super-sphere shot", says Okun. "That starts in the swamp and takes you to various Spheres activating around the world. That took 30 days to render. That's pretty crazy. It's around 1,100 frames. It's an amazing shot. You don't want to show it to the director at the end of the day and have him say, 'That's not really our sphere'...which is what happened. We came up with a patch system at Weta Digital where we could render a section and patch it over the offending thing. This particular patch took three days to render."

Weta Digital

"Linux is an integral part of what we do here at Weta", says Production Engineering Lead Peter Capelluto. "It's very well suited for the dynamic needs of the visual-effects industry. Our department would have a much more difficult time accomplishing our goals with any other operating system."

"Weta predominantly uses Linux for our workstations and also for our renderfarm and servers", says Capelluto. "There are a few applications that require the use of Mac OS X, Windows and Irix. Whenever possible, we use Linux. The open-source nature of Linux and the many Linux applications are a major advantage. We also prefer it for stability, low cost, access control, multiuser capabilities, control and flexibility." Capelluto's department develops pipeline software, such as the digital asset management system and the distributed resource management system for their renderfarm.

"We have 500 IBM Blade Servers, 2,560 HP BL2x220C Blade Servers and 1,000 workstations", says Weta Digital Systems Department Lead Adam Shand. "Ubuntu is our primary render and desktop distro. We also use CentOS, RHEL and Debian." The workstations are IBM and HP. Weta uses NetApp DataOnTap, NetApp GX, BluArc, Panasas and SGI file servers. Storage is mostly NAS, not SAN. For open-source apps, they use Apache, Perl, Python, MySQL, PostgresSQL, Bind, OpenOffice.org, CUPS, OpenLDAP, Samba, Firefox, Thunderbird, Django, Cacti, Cricket, MRTG and Sun Gridware.

"We're big fans of open-source code here at Weta", says Capelluto. "We're utilizing Sun's Grid Engine for distributed resource management and have helped them fix a number of bugs. It's very powerful to be able to improve upon open-source software and to fix any problems you encounter."

Cinesite

"When your supervisor is in New Zealand and your editor is in Los Angeles, it makes things a bit harder", says Cinesite Visual Effects Producer Ken Dailey, who is based in London. "I'd talk to Jeff every day and make sure he has the right Quicktimes, that everyone is looking at the same stuff. Time was the biggest challenge. The reaper shot in New York came very late. I think we had three weeks from the time we got the plates. We shared Maya models with Weta. We sent them our reaper model and they shared models with us."

"We did about 60 shots", says Dailey. "We did where Klaatu is being interrogated, which shows how he can take control of electrical systems. Later in the movie, there's a sequence where the military decides to attack Gort in Central Park with drones. We had 3-D tanks and explosions. We did the big tilt-down from space at the beginning of the movie."

"We're principally using Maya, Shake and RenderMan", says Dailey. "Shake is running on Linux. Maya is running mostly on Windows. We use a bit of Photoshop on Windows." Cinesite uses Maya on Linux and Windows. However, the range of available plugins is far greater on Windows. The 3-D painting package MudBox is the main one. That's recently been bought by Autodesk and may be coming to Linux.

"We have about 80 desktop systems running Linux", says Cinesite Senior Systems Administrator Danny Smith. "We have at least a couple hundred render systems. All of those are IBM Blade systems. We have about 40 Windows systems as well. Our principal requirement for Windows is Photoshop. There's no way to run Photoshop reliably in its latest release on Linux with CrossOver. The main reason for Photoshop is the color depth—the full 16 bits we require for film work in matte painting and dealing with textures."

"CinePaint was looked at in the past, back when it was Film Gimp", says Smith. "Our biggest problem with it is finding people with the skills to use it. People walking in the door already know Photoshop. People would be more interested in GIMP or CinePaint if it was more like Photoshop. If we

could find a tool that reduces our Photoshop costs, a lot of people would be very happy. We have 20 or 30 seats of Photoshop."

"Shake is a product that's being discontinued", says Smith. "Even though we've done the site buyout, as soon as Apple launches a competing product, they have the right to discontinue our use of Shake. The likely successor is Nuke. We're trying to get people up to speed with Nuke and doing more and more with it. It takes time to train people. It's slowed down our adoption."

"We mostly run Red Hat Fedora", says Smith. "We're on version 4, migrating to 8. We've experimented with SUSE. The reason to stay with Red Hat is support from software vendors. We're paying for that support, and it's mission-critical."

For dailies playback, Cinesite is using the Windows system Scratch from Assimilate. Scratch also is being used by the Avid editors in Los Angeles on the Fox lot. Smith had the Linux SpecSoft RaveHD dailies system at his prior company, but considers the California startup too far away to support London. Cinesite also uses FrameCycler on Linux for movie playback. They have NetApps and Isilon file servers.

Flash Film Works

"We had Flash Filmworks handle a hundred shots, 3-D helicopters and stuff like that", says Okun. Flash Film Works, based in Los Angeles, has its desktops set up to dual-boot. "This was one of the rare occasions where most of the workstations stayed in Windows", says Flash Film Works Technology Chief Dan Novy. "That's mostly because we weren't doing a lot of fluid dynamics simulations. The renders were 80% Windows. I didn't need the high performance that I normally use 64-bit Fedora for. The file servers are all Linux. I have a specialized Shake station that has The Foundry Furnace suite on it for doing automation."

Even running Windows, they still are using open-source tools. "Fusion 5 added Python in addition to its Lua-based Ion scripting", says Novy. "That can do a lot more automation, getting renders to the editor automatically, that sort of thing." Fusion recently has become available on Linux, but Lightwave is Windows. Flash Film Works likes Lightwave's free render nodes. A Maya RenderMan node would cost them $4,000–$5,000; Mental Ray costs $1,200.

"I personally run Ubuntu on my laptops", says Novy. "But, for setting up file servers, I'm so used to the Red Hat paradigm. We have one Isilon cluster, a FreeBSD variant. Each node is 1.4TB. I have five nodes and one backup. It's old, and I'm leaning toward BluArc to replace it. We have about 100 CPUs on the farm and about 50 desktops." Flash Film Works backs up its data to Blu-ray.

The arrival in New York's Central Park of a giant sphere from another world may have dire consequences for our planet. Photo credit: WETA.

Klaatu's (Keanu Reeves) arrival on Earth via a giant sphere triggers a global upheaval. Photo credit: WETA.

Hammerhead

"Hammerhead did a handful of really difficult and excellent shots", says Okun. "They did one really difficult 3-D helicopter with the Secretary of State on it. They did a matte painting that was really 'saving the day' that had an automobile factory. They did the trooper-healing shot."

"Hammerhead uses Linux for all of our graphics workstations for our visual-effects artists, as well as for our render boxes, and file servers", says Hammerhead Visual Effects Supervisor Thomas Dadras. "I feel that Linux is the best possible environment for visual-effects production because it's so incredibly customizable and scalable. We utilize a full spectrum of in-house scripts, aliases and environment variables that enable artists to easily navigate the file systems for the many shows that we have in-house at any given time."

A massive and seemingly unstoppable alien force destroys a stadium in seconds. Photo credit: WETA.

Hammerhead uses Maya with RenderMan and Mental Ray for rendering, Shake and Nuke for 2-D compositing and rotoscoping, Photoshop for texture painting and matte painting, and SynthEyes for 3-D tracking. They've internally developed software for 2-D and 3-D tracking and rotoscoping.

"The company size fluctuates between ten to 25 people depending on the amount of work", says Dadras. "We currently have 17 people. We have nine artist desktop-Linux workstations, all with dual monitors. The capacity of our Isilon server is 17TB. We also have an older SGI file server, about 5TB in size. We currently have 22 render blades on our renderfarm."

"All our workstations and render nodes are running on CentOS 5", says Hammerhead Systems Administrator Fatima Mojaddidy. Hammerhead has eight Macs that are used by producers and coordinators for running software such as Filemaker. A few of the Macs also are used in production with Photoshop and SynthEyes.

Hydraulx

"Hydraulx stepped in right near the end", says Okun. "They did the most amazing job of creating an army out of eight jeeps and 50 people. They were the only other people on the show to deal with any particle systems stuff."

"All our workstations, our entire facility is now Linux", says Hydraulx Visual Effects Supervisor Colin Strause. "We use Inferno and Flame, Shake, Maya, Photoshop, a little bit of After Effects, Combustion, Synflex for cloth simulations, Real Flow for fluid stuff and Massive. Everything after that is our own custom tools. We use GIMP for doing mid-level painting stuff—quick texture stuff. We've gone dual-boot on most of our Linux machines, so modelers and texture modelers can run Z-brush and Photoshop."

"We have more than 500TB on the SAN network", says Strause. "We have a Think Logical KVM switch, based on fibre, that will route your monitor, keyboard, mouse and tablet. You can take any machine in the building—our 25 Inferno stations or three big Final Cut Bays—and route it to any other machine in the building or up into the screening room." (The screening room has a 23'-wide screen.)

"Our sequence takes place where the Swarm escapes from a missile silo where they were storing Gort", says Strause. "The US military has hundreds of tanks and soldiers and missile launchers there in case something bad happens. The Swarm takes out the whole army. We were able to rent a handful of vehicles...some M-1 Abrams, some Bradley Fighting Vehicles, a bunch of Hummers with 50-caliber machine guns on the top and troop transport trucks."

"We had only six to seven weeks to do the entire sequence", says Strause. "Normally, you'd have three or four months. The other problem we had was matching all the Swarm dynamics that Weta did. Each company has such a unique pipeline, there's very little we can share. We looked at the trailer from Apple.com to figure it out and reverse-engineer their Swarm effects to create them from scratch for our shot."

"With all the fires in Los Angeles, we couldn't shoot the weapons", says Strause. "So for all the weapons you see firing, we had to add CG shell casings, and we created all the tracers and muzzle flashes with fluid simulations."

"When a 60-ton tank shoots, it's going to shake the ground, so all the dust comes into the air", says Strause. "When the tanks fire, we have all the correct dynamics, such as the heavy tank tread jiggling on the suspension. We went through YouTube, which is great, and found all these videos that guys had taken in Iraq of their tanks shooting. It was an amazing reference."

"On set, we have a Sphere-On camera that lets you take 360-degree HDRI images", says Strause. "We use these super-high dynamic images for photometric lighting, trying to re-create how the real light behaves in the real world in our digital environment. When we have a real tank and a CG tank right next to it, we have to use a much fancier technique to make it all look photo-real."

Hydraulx photographed the vehicles at a desert location in Los Angeles, then modeled everything in Maya. Camera tracking uses Linux Boujou software, brought into Mental Ray for lighting and shading.

"We use digital crowd simulation software called Massive", says Strause. "We've written some custom tools that let us get the stuff into Maya so we can render it in our Mental Ray pipeline. We have soldiers with guns and they're running. The soldiers avoid the moving vehicles automatically. It's all done with this neural network."

"We have a custom version of Piranha here that we use for all our dailies", says Strause. "We have it on every single Linux machine. We have an elaborate database, based on MySQL, that's our shot management and render manager. We can dynamically build edits. We can take an EDL of an off-line, and whenever people mark a new daily that they want to look at in the cut, they can hit View and Cut, and it will dynamically build off an XML file all the different QuickTimes, load Piranha and dynamically build the cut. We use Shake to build the QuickTimes as a job on our Linux renderfarm."

Conclusion

"Linux is the OS of choice for the film industry", says Autodesk Television Industry Manager Bruno Sargeant. "Linux leverages the processing power of current workstations and allows Flame and Inferno to reach performance previously obtainable only with supercomputers like SGI."

Even more important than Linux are the artists using Linux. "Computers are machines, and they're no different from a paintbrush", says Okun. "It's the artist running the computer who makes the difference."■

Robin Rowe is a Linux consultant and the project leader of CinePaint (www.CinePaint.org).

Show Highlights

- Security
- Enterprise Systems
- Linux Installations
- Open Source Solutions
- Business Solutions
- Education Solutions
- Linux Integration
- Linux Certification by Linux Professional Institute (LPI)

March 9, 2009
8 AM - 5 PM
University of North Florida University Center
Jacksonville, Florida

http://www.floridalinuxshow.com/

Donald L. Corbet "TheLinuxGuy®"
2009 Show Chairman

Rod Sharp
2009 Show Director

Keynote Speaker:
Mr. Sam Carpenter

Linux Overlord Presenter:
Mr. Joe "Zonker" Brockmeier

Special Guest:
Mr. Robin "Roblimo" Miller

Charitable Outreach

Children's Miracle Network®

We believe it's important to support worthwhile charitable causes. This year, Florida Linux Show will donate $1.00 of every $25.00 show admission to the Children's Miracle Network. Donation jars will also be available during the show for anyone who wishes to make an additional contribution.

Official Sponsors

GOLD

GSA PROPOSAL.COM

SILVER

O'REILLY®

ThinkGeek

BRONZE

Zenoss®

MEDIA

LINUX JOURNAL

LinuxQuestions.org

LINUX PRO MAGAZINE

JOB LISTINGS

COMPUTERWORK.com™

Supporting Florida LUGs: GatorLUG / Linux enthusiasts & professionals / Pensacola Linux User's Group / Florida Linux Users Group / Linux Users Group of Tallahassee/ Thomasville / Jacksonville Linux User's Group

Florida Ubuntu Loco Team

OSS Atlanta.org

Open Source Software

Run Your Own Virtual Reality with OpenSim

By running OpenSim on your own server, you can create your own virtual world and even connect it to other virtual worlds.

BILL CHILDERS

Last month, I took you on a tour of *Second Life* and showed you some applications for it besides entertainment. This month, I explain how you can run your own virtual island using the *Second Life* client and the open-source 3-D application server software called OpenSimulator. OpenSimulator, or OpenSim for short, is a free (as in speech) implementation of a virtual world platform, utilizing the *Second Life* protocols. From its FAQ: "OpenSim is a platform for operating a virtual world and supports multiple independent regions connecting to a single centralized grid. This is somewhat similar to the Web, where people can run their own Web servers, tied together through the Internet. It can also be used to create a private grid, analogous to a private intranet." In other words, OpenSim can be like a 3-D Apache server, enabling collaboration, entertainment and business without having to utilize a centralized service.

You don't need a beefy machine to run OpenSim. I've had a Pentium 4 with 1GB of RAM running a basic island for a little while, and it has supported eight concurrent logins without any trouble. If you have a more modern machine, you even can run OpenSim locally on the same machine as your *Second Life* client. I also have tested this configuration on my Core 2 Duo-equipped laptop, and it ran quite well.

The client is a much more demanding system than the server is. I hardly notice the server process when running this in "all-in-one" mode. As a matter of fact, one machine can support a number of sims (also known as regions). I have tested up to ten regions on a dual-core machine with 2GB of RAM. So long as there isn't crazy activity in one of the regions, they all have acceptable performance. This is similar to what Linden Labs does in its data centers—it runs its "openspace" regions at four regions per CPU core.

The folks over at OpenSim say that the best method for getting and running OpenSim is to build it from source, though they do offer nightly builds. OpenSim is written in C# and runs under Mono on Linux. Due to the high dependency on Mono, I'm assuming you are running Ubuntu 8.04 or 8.10 (to get OpenSim running on other distros, please see the Build Instructions in the Resources section). Versions of Ubuntu earlier than 8.04 need to have Mono upgraded at least to version 1.9.1. (OpenSim will run under the 1.2.6 that ships with Ubuntu 8.04, but 1.9.1 is preferred for its better memory management.) The build process is well documented and straightforward, even for a ham-fisted guy like me.

Before you try it, please check the OpenSimulator Build Instructions link in the Resources section of this article, but here's how you build OpenSim on Ubuntu 8.04. This set of commands installs Mono and the necessary dependencies and libraries, gets the latest source from SVN and builds it:

```
sudo apt-get install subversion nant mono-gmcs \
             libmono-microsoft8.0-cil \
             libmono-system-runtime2.0-cil \
             libgdiplus libmono-i18n2.0-cil ruby
svn co http://opensimulator.org/svn/opensim/trunk opensim
cd opensim
./runprebuild.sh
nant
```

Once nant finishes running, that's it. OpenSim is ready to be configured and started. OpenSim can run in two modes: standalone and grid. Standalone OpenSim servers are islands unto themselves—self-contained and not dependent on any separate login managers or services. Grid mode is the more interesting mode, where multiple islands can be linked together into a cohesive virtual world. *Second Life* itself is an example of a bunch of regions running in grid mode, although there are other third-party grids now, such as OSGrid, DeepGrid and Openlife Grid.

Before you configure your OpenSim, you need to decide on the mode in which you will run. I recommend running your first sim in standalone mode to get the hang of the way things work. Once you understand standalone mode, you can set up another sim or convert your standalone sim to a grid-based sim, and either build your own network of sims or connect your sim to another grid altogether, like OSGrid. Configuring a grid server is beyond the scope of this article, but there is excellent documentation at the OpenSimulator Web site.

Configuring a basic standalone server couldn't be easier, as the server asks you several questions on first startup, and it builds the opensim.ini configuration file from that.

Another decision you need to make prior to configuring OpenSim is what database back end you want to run. For experimentation and light use in standalone mode, the default SQLite database is perfect. There are no special configuration settings to mess with and no external databases to set up. When you decide to graduate to grid mode, however, you probably will want to run the back end on MySQL. Configuring a MySQL back end also is beyond the scope of this article, but again, there is great documentation on this at the OpenSim Web site.

Let's get your first OpenSim up and running. All the

compiled binaries will be in the opensim/bin directory. You either can run OpenSim from there or move them to a directory of your choice. If you choose to move them, ensure that you move the files under the opensim/bin directory recursively. Once you've got the binaries in place, you can start OpenSim simply by running `mono ./OpenSim.exe` from the directory in which you put the binaries. Once that command is executed, OpenSim goes to work, creating a bunch of standard "assets" and populating the SQLite database.

Next, you will see a prompt that asks you about the configuration of your server. It's actually safe to accept all the defaults if you just want to play with an OpenSim on the same machine as the client. Note that OpenSim's default is to listen on loopback (that is, localhost) only. If you don't put the IP address of your machine in the external hostname field, you won't be able to connect to your OpenSim instance from another machine. You definitely will need to deviate from the defaults if you want to have the client and server on different machines. You probably also will want to specify the Master Avatar first and last name as something other than the default. Here's an example:

```
21:45:34 - [CONFIG]: Calling Configuration Load Function...
DEFAULT REGION CONFIG: Region Name [OpenSim Test]:
Bill's Place
DEFAULT REGION CONFIG: Grid Location (X Axis) [1000]:
1000
DEFAULT REGION CONFIG: Grid Location (Y Axis) [1000]:
1000
DEFAULT REGION CONFIG: Internal IP Address for
        incoming UDP client connections [0.0.0.0]:
0.0.0.0
DEFAULT REGION CONFIG: Internal IP Port for
        incoming UDP client connections [9000]:
9000
DEFAULT REGION CONFIG: External Host Name [127.0.0.1]:
192.168.1.140
DEFAULT REGION CONFIG: First Name of Master Avatar [Test]:
Bill
DEFAULT REGION CONFIG: Last Name of Master Avatar [User]:
Deere
DEFAULT REGION CONFIG: (Sandbox Mode Only)Password for
        Master Avatar account [test]:
password
```

Once you answer these questions, the OpenSim server finishes its startup sequence and is ready for you to log in when you see a `Region (root) #` : prompt. The OpenSim console allows you to do some basic management and administration of the server from this prompt, but the real fun stuff begins when you fire up a *Second Life* viewer. The actual virtual world can be accessed only via the *Second Life* viewer. Assuming you already have the *Second Life* viewer and have untarred it into a directory, you can log in to your server by running the following command:

```
./secondlife -loginuri http://<server IP>:9000/ \
             -loginpage http://<server IP>:9000/?method=login
```

In my case, I can log in to my server with the following:

```
./secondlife -loginuri http://192.168.1.140:9000/ \
             -loginpage http://192.168.1.140:9000/?method=login
```

The client will launch, and you will see a basic HTML screen (Figure 1). Input the first and last name of the Master Avatar that you specified when configuring the server (see above) and the password you specified.

The client then will log in to the server, and you'll be on your own little island! It's not much to look at, just a little hump in an empty sea, but then again, neither is your avatar. Regardless of how you envision your avatar, all avatars in an OpenSim deployment start out as *one* basic shape—the basic Linden shape known as Ruth. This is because you have no prepopulated inventory, so you get the generic unisex avatar, which just happens to look female. Go figure.

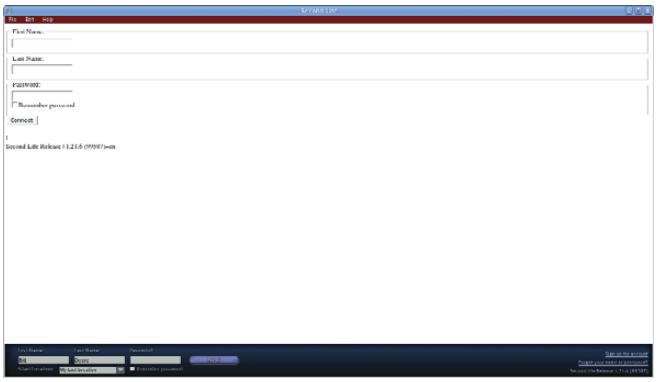

Figure 1. Logging In to My Virtual World

Ruth is sort of ambiguous-looking (Figure 2), but it's not hard to fix that. By going into the inventory window, right-clicking on Body Parts, and then clicking on New Body Part—New Shape, you can create a body shape, then wear it. Once you wear that shape, you can right-click on your avatar, and select Edit Appearance. Then, edit that shape to your liking. Once you get your shape dialed in, you can create basic clothing that's a little better than what's provided with Ruth, and look as fashionable as I do (Figure 3).

Another thing you may want to do to your island is to terraform it. This is done by right-clicking on the ground and selecting Edit Terrain. The terraforming dialog comes up, and from there, you can use the raise land or lower land tools to change the elevation of certain areas of the parcel. I spent about five minutes in the terraforming dialog, and my island went from the basic little round hump to something that vaguely resembles Pac-Man (Figure 4).

If you are a seasoned *Second Life* user, you may go through a period of shock once you realize there is no prepopulated inventory, nor are there any stores to buy goods on your own private island. Much like Robinson Crusoe, you are going to have to make anything you want to use or wear on your own

Figure 2. What I Looked Like at First

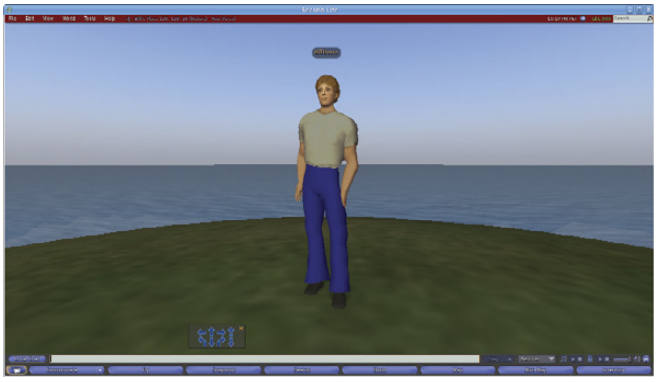

Figure 3. What I Looked Like after a Makeover

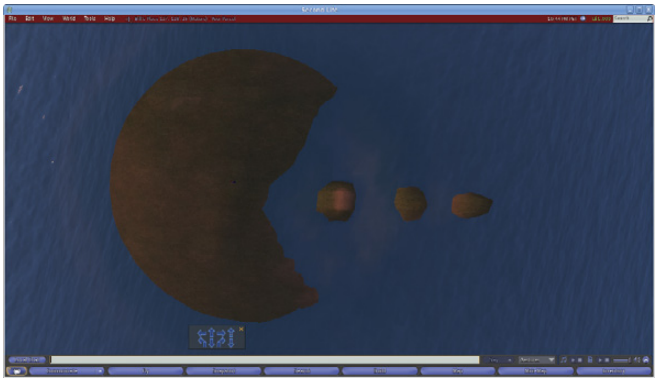

Figure 4. My Island after Its Makeover

island. If you do know LSL scripting in *Second Life*, you'll be glad to know that there is an analog in OpenSimulator as well. The OpenSim folks have implemented about 85% of the LSL scripting language in OpenSim, and they are improving the support for it every day.

Once your island is up and running, for an added challenge, you may want to try connecting your private island to one of the alternative grids that have been set up by virtual reality fans. See the Alternative Grids sidebar for more information on this—it's a great way to jumpstart into a whole new virtual universe. If you can set up the infrastructure, but have a difficult time creating content, you may choose this path.

Once you've hooked your island into a grid, you can start to leverage all the shared inventory and assets of other people who have come before you. Much like *Second Life*, there are stores on the alternative grids where you can get clothes,

Alternative Grids

Some of the alternative grids are very interesting places to visit if you're adventurous. One of my favorite places to check out is OSGrid. OSGrid is a loose confederation of independent OpenSim operators who have tied their sims into a cohesive grid. Some of the operators are universities experimenting with virtual worlds, and others are amateur enthusiasts or OpenSim developers. Getting onto OSGrid is easy; simply go to the OSGrid site (see Resources), and register for a free account. Then, start your *Second Life* client with the following command:

```
./secondlife -multiple -loginuri http://osgrid.org:8002 \
             -loginpage http://osgrid.org/loginscreen.php \
             -helperuri http://osgrid.org
```

Once you log in, feel free to wander about (Figure 5). There's lots to see and do, and unlike *Second Life*, there is no economy, so you either find and utilize freebie items or you create your own content. OSGrid is a great sandbox for builders in *Second Life* who want to experiment without upload costs as well.

If you want to try running your own OSGrid-enabled server at home, you can find instructions on how to join your own server to OSGrid at the OSGrid Web site. Attaching your own server to OSGrid is much simpler than running your own private grid, because the centralized login and asset servers are maintained by the OSGrid group. Reserve a spot on the OSGrid map for your sim and check it out! It's easier than it looks.

Figure 5. Me Stepping Out to OSGrid

Getting Help

You can get help with your own OpenSim deployment via various methods. Possibly the most convenient method is the OpenSim IRC server that you can log in to for help with running your own server. There is a great FAQ on the OpenSim site as well. Other great resources are the residents of OSGrid. Because they all had to connect their sims to OSGrid themselves, they can be great technical resources.

vehicles, furniture and other goods—and most of it on the alternative grids is free!

That's the basics on your own island. At this point, you can have friends or colleagues log in to your server using the *Second Life* client, and you can converse and collaborate. Explore what you can do with the technology, and have fun making your own virtual world!■

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. If he ever finds the time to make it into *Second Life*, he goes by Bill Deere.

Resources

OpenSimulator: **www.opensimulator.org**

OpenSimulator Build Instructions: **opensimulator.org/wiki/Build_Instructions**

OpenSimulator Configuration Instructions: **opensimulator.org/wiki/Configuration**

Connecting to OpenSim: **opensimulator.org/wiki/Connecting**

OpenSim Scripting: **opensimulator.org/wiki/Scripting_Documentation**

OpenSim IRC Channel: **irc://irc.freenode.net/opensim** (#opensim on irc.freenode.net)

Video: How to Run OpenSim on Windows: **blip.tv/file/1421954**

Second Life Viewer Download: **secondlife.com/support/downloads.php**

OSGrid: **www.osgrid.org**

DeepGrid: **www.deepgrid.org**

Openlife Grid: **openlifegrid.com**

ASA COMPUTERS

Want your business to be more productive?
The ASA Servers powered by the Intel Xeon Processor provide the quality and dependability to keep up with your growing business.

Hardware Systems for the Open Source Community - Since 1989.
(Linux, FreeBSD, NetBSD, OpenBSD, Solaris, MS, etc.

1U Server - ASA1401i
- 1TB Storage Installed. Max – 3TB.
- Intel Dual core 5030 CPU (Qty-1), Max–2 CPUs
- 1GB 667MGZ FBDIMMs Installed.
- Supports 16GB FBDIMM.
- 4X250GB htswap SATA-II Drives Installed.
- 4 port SATA-II RAID controller.
- 2X10/100/1000 LAN onboard.

2U Server - ASA2121i
- 4TB Storage Installed. Max – 12TB.
- Intel Dual core 5050 CPU.
- 1GB 667MGZ FBDIMMs Installed.
- Supports 16GB FBDIMM.
- 16 port SATA-II RAID controller.
- 16X250GB htswap SATA-II Drives Installed.
- 2X10/100/1000 LAN onboard.
- 800w Red PS.

3U Server - ASA3161i
- 4TB Storage Installed. Max – 12TB.
- Intel Dual core 5050 CPU.
- 1GB 667MGZ FBDIMMs Installed.
- Supports 16GB FBDIMM.
- 16 port SATA-II RAID controller.
- 16X250GB htswap SATA-II Drives Installed.
- 2X10/100/1000 LAN onboard.
- 800w Red PS.

5U Server - ASA5241i
- 6TB Storage Installed. Max – 18TB.
- Intel Dual core 5050 CPU,
- 4GB 667MGZ FBDIMMs Installed.
- Supports 16GB FBDIMM.
- 24X250GB htswap SATA-II Drives Installed.
- 24 port SATA-II RAID. CARD/BBU.
- 2X10/100/1000 LAN onboard.
- 930w Red PS.

8U Server - ASA8421i

- 10TB Storage Installed. Max – 30TB.
- Intel Dual core 5050 CPU.
- Quantity 42 Installed.
- 1GB 667MGZ FBDIMMs.
- Supports 32GB FBDIMM.
- 40X250GB htswap SATA-II Drives Installed.
- 2X12 Port SATA-II Multilane RAID controller.
- 1X16 Port SATA-II Multilane RAID controller.
- 2X10/100/1000 LAN onboard.
- 1300 W Red Ps.

All systems installed and tested with user's choice of Linux distribution (free). ASA Collocation—$75 per month

2354 Calle Del Mundo,
Santa Clara, CA 95054
www.asacomputers.com
Email: sales@asacomputers.com
P: 1-800-REAL-PCS | FAX: 408-654-2910

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

Prices and availability subject to change without notice. Not responsible for typographic errors.

intel Xeon® inside™
Powerful. Efficient.

EOF

The Most Personal Device

Today it's the iPhone. Tomorrow it's a Linux handheld.

DOC SEARLS

Several years ago, I gave my wife a Nokia phone that I hoped would get her to start texting. It was a small phone that twisted open to expose a qwerty keyboard. I showed her how to use it, and later sent her a test message. Her reply was "no". She never texted again with that or any other phone.

Until last December—that's when she got an iPhone. Within days, she was expert at texting and is still schooling me on how to use the iPhone I bought for myself last summer.

She calls the iPhone her "laptop replacement". She still uses a laptop, but its main job is to serve as a wide-screen iPhone that also syncs calendars, apps and podcasts.

Yes, I know the iPhone is not a Linux device. But that's not the point. The iPhone is the modern equivalent of the Apple II. It models the future and cracks open a vast new territory for development.

The iPhone is the first phone that subordinates telephony to the rest of what it does, which could be anything. It's a handheld computing device that happens to do telephony. It's native to the Net, not just to the phone system. And it opens a category that Android and other Linux-based phones will fill. That category will have at least these three virtues: 1) it will be native to the Net, not just the phone system; 2) it will be generative—that is, it will open rather than closed to the possibilities for what can be developed for it; and 3) it will expand the range of what individual human beings can do while moving about in the world.

That last virtue is not shared with desktops or laptops, because those are mostly limited to what you can do sitting down.

At LinuxWorld August 2008, I went by the Access booth (**www.access-company.com**), where they were showing off the vast range of devices using the company's Linux development and deployment systems. These included mobile phones, PDAs, Internet terminals, car navigation systems, set-top boxes, business operation terminals, musical instruments, video game consoles, IP phones, home appliances and other devices. But mobile phones were the main thing. Amid Access' literature was a poster showing off 219 different "Access-powered mobile phones". All running on Linux, presumably.

A couple Access employees greeted me, and I hit them right away with question begged by the iPhone's success: "How long before the cell phone companies realize they're running a data system and not a phone system?" They were taken aback at first, but gave thoughtful responses. "It'll be a long time", one guy said. "But it will happen", the other guy said. (See "WiMAXing Linux" on page 16 for more on that.)

Then we started talking about the mobile data business, which in their case was supporting development of apps for "Access-powered" phones. I asked if their system supported audio yet. One of the guys said no. At this point, I felt comfortable pulling out my iPhone and showing one app among many that was changing my life: a stream tuner for Internet radio. It wasn't perfect, but because of it, my iPhone had become my main radio. I can "tune" in .mp3 streams from anywhere that has an exposed URL or IP address. I can listen anywhere in the US for however long I like. In cars, I jack it into the AUX input on the dashboard. Thanks to the unlimited data deal I have with AT&T, I don't worry about drinking too many bits.

After showing the two guys my iPhone playing a Boston radio station, both of them felt comfortable pulling iPhones out of their pockets as well.

My friend Keith Hopper made an interesting observation recently. He said one of Apple's roles in the world is finding categories where progress is logjammed, and opening things up by coming out with a single solution that takes care of everything, from the bottom to the top. Apple did it with graphical computing, with .mp3 players, with on-line music sales and now with smartphones. In each case, it opens up whole new territories that can then be settled and expanded by other products, services and companies. Yes, it's closed and controlling and the rest of it. But what matters is the new markets that open up.

Android phones began hitting the streets late last year. They aren't as slick and easy to use as the iPhone, but that doesn't matter. In two years, all current models of both will be very old hat. What matters is that Android is Linux-based and an open platform. Those two facts alone will help accelerate the inevitable conversion of the cell-phone system to the cell-data system.

Android and other open platforms won't just be media recorders and players, game machines, phones, musical instruments, radios and texting devices. They will become wallets. They will shake hands for us and help us do business. They will help us be more of what we are, which is human.

As creatures, we humans are distinguished not only by our intelligence and use of language, but also by two other remarkable characteristics: our mobility and our expansiveness. We are relatively hairless and walk on two feet because we are runners. A well-conditioned adult human can run indefinitely. We also expand our very selves though the things we invent, hold and manipulate. Our senses spread out through our clothes, our tools and our tech by a process called indwelling. When drivers say "my wheels" or pilots say "my wings", they mean it personally. The perimeters of our selves are not bound by our bodies. They extend to include the tech we use. To become expert is to enlarge ourselves, whether as carpenters, drivers, pilots or whatever.

There is an evolutionary progression from desk to lap to palm. Apple has done us the favor of pointing the way. Our job is to follow the path and open the territory. When we're done, "desktop" and "laptop" will sound as antique as "mainframe" and "minicomputer".■

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.

Rackspace® Hosting

Proudly supports our most important partner

Mother Nature.

Welcome to Greenspace™

Now, you can make your company greener just by hosting at Rackspace.

With our Greenspace initiative, hosted servers are both business friendly and earth friendly.

- Energy efficient components so they run even greener
- Same cost and performance as regular servers
- Choose from standard green configurations or create your own

Fanatical Support® has always been about doing right by our customers.
Now it's about doing right by Mother Nature.

Microsoft redhat DELL CISCO hp AMD vmware PCI

rackspace® HOSTING

experience ***fanatical support***®

rackspace.com/linuxjournal • 888-571-8976

GPU Computing

More GFLOPS, Fewer Watts!

Clusters with Tesla S1070 - 4 GPU Servers

- 36 GPUs + 36 CPUs + 24 TB in 24U
- 40 Gbps FasTree™ InfiniBand
- InfiniScope™ Network Monitoring

WhisperStation™ with 1 to 4 Tesla GPUs

Tesla C1060 GPU Performance:

- 1 TFLOPS per GPU
- 4 GB DDR3 per GPU
- 102 GB/Sec Bandwidth

FireStream™ Workstations & Servers

WhisperStation™ – FS with FireStream 9250 Stream Processors

AMD FireStream 9250

Stream Processor Performance

- 1 TFLOPS
- 1 GB DDR3
- Single Wide Card /150 Watts
- Double Precision Floating Point Hardware

Stream SDK

- Brook+ Open Source Compiler
- ACML and CAL Math Libraries

Microway®
Technology you can count on℠

508-746-7341
microway.com

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