AN SSC PUBLICATION

PROFESSIONAL DVDS FROM HOME VIDEO USING QDVDAUTHOR

LINUX JOURNAL™

Since 1994: The Original Magazine of the Linux Community
DECEMBER 2005

amaroK
The music manager for Linux that Rocks!

Ruby on Rails
Working with ActiveRecord

mythTV is no MYTH!

RECORD YOUR FAVORITE TV SHOWS FOR FREE WITH THIS LINUX-BASED DIGITAL VIDEO RECORDER

NEW COLUMN ON LINUX SHELL SCRIPTING: "WORK THE SHELL"

Secure Your Corporate Directory with OpenLDAP and OpenSSL

Stunning Graphics with GNU Lilypond

tutorial: Using vmstat for Tuning Virtual Memory

USA $5.00 CAN $6.50
www.linuxjournal.com
12
0 71486 03102 4

LINUX JOURNAL · MULTIMEDIA · MythTV · Ruby on Rails · amaroK · OpenLDAP · vmstat · GNU Lilypond · DECEMBER 2005 ISSUE 140

THE UNSUNG HERO:

Network Administrator

9:42 am Singapore branches go offline, trouble ticket created

9:44 am Jeff uses diagnostics to isolate failure to core router – not responding

9:45 am Out-of-band access to core router established via the **AlterPath™ ACS**

9:47 am Router shows subnet mask set incorrectly during previous configuration

9:48 am Jeff resets subnet mask properly, reboots router

9:49 am Link to Singapore restored, Singapore comes back online

9:50 am Jeff is planning his next vacation

Advanced Console Server

AlterPath™ ACS

Download a **FREE** White Paper on Console Management
www.cyclades.com/wpcm

www.cyclades.com/ljb
1.888.cyclades · sales@cyclades.com

cyclades

©2005 Cyclades Corporation. All rights reserved. All other trademarks and product images are property of their respective owners. Product information subject to change without notice.

The Apache Software Foundation and Full Circle Productions present

ApacheCon US '05

Leading the Wave of Open Source

December 10-14, 2005 • San Diego, CA

Tutorials • Conference • Expo • Networking • Special Events

Hone Your Competitive Edge

Learn shortcuts, tricks, and techniques that today's developers need to succeed. More than 100 sessions highlight core and next-generation Apache server issues aimed at both the technologist as well as implementer, including:

Apache HTTP Server • Derby • Geronimo • Lenya • James • Axis
Lucene • Cocoon • Tomcat • SpamAssassin • Tapestry • Struts
Forrest • Maven • PHP • Perl • Subversion . . . and more!

Foto: Thor Søndeaa

"We are witnessing a shift of power in the growing influence of Open Source, bringing a critical need for a new generation of developers to provide innovative solutions at all levels. We are committed to help educate and nurture the community by sharing ideas and experiences with leading industry experts, as well as independent developers, academics, and visionaries."
— **Ken Coar**, ApacheCon Chairman

Special Events Include:

- Receptions
- Birds of a Feather (BOF) Sessions
- ApacheCon Labs
- PGP Key Signings
- Help Desks
- Book Signings
- Lighting Talks
- The World Premiere Screening of **"FUD"**, the documentary on Open Source filmed at OSCOM 4 and ApacheCon US 2004

Premier Sponsors Include:

Sun microsystems · BlackHat · Java Community Process · covalent · Google · IBM

YAFIYGI • LLTA

Cory Doctorow and **Jaron Lanier** headline a top-tier roster of speakers and presenters, including Brian W. Fitzpatrick, Rasmus Lerdorf, Ted Leung, Geir Magnusson Jr., Geoffrey Young, Philippe M. Chiasson, Paul Querna and Cliff Schmidt.

Open for Business

Simon Phipps, Sun Microsystems Chief Open Source Officer, kicks off a one-day Open Source and Business Track with "The Zen of Open".

Other Sponsors and Exhibitors Include:
Azul Systems, cafésoft, Rogue Wave Software and thawte

Media Partners Include:
BYTE.com, C/C++ Users Journal, developer pipeline, Dr. Dobb's Journal, Eclipse, InfoWorld, Java Developers' Journal, Linux World Magazine, O'Reilly Network, PowerBuilder, Security Horizon, Software Development Online, Sys Admin, The Perl Journal, Unix Review, Web Services Journal, and XML Journal

Register NOW at http://www.apachecon.com

COVER STORY

MYTHTV IS NO MYTH

In this month's *LJ*, James Turner provides an overview of MythTV, a Linux-based TiVo replacement (page 64). And, Matthew Gast, in his "Advanced MythTV Video Processing", shows you how to deinterlace video playback and extract video so you can take recorded programs on the road (page 69).

LINUX JOURNAL

DECEMBER 2005 ISSUE 140

FEATURES

48 MAKE STUNNING SCHENKER GRAPHS WITH GNU LILYPOND
Think that Linux-based graphics programs can't compete? GNU Lilypond generates beautiful graphics that make commercial alternatives seem second-rate.
KRIS SHAFFER

56 DVD MASTERING USING QDVDAUTHOR
Do you need a Mac to do professional DVD authoring? Not at all. It's easy with QDVDAuthor.
DAN SAWYER

64 A LINUX DVR IS NO MYTH—IT'S MYTHTV!
Although commercial Digital Video Recorders (like TiVo) charge more and restrict your options, Linux alternatives are free and rock!
JAMES TURNER

INDEPTH

69 ADVANCED MYTHTV VIDEO PROCESSING
Once you've got MythTV set up, how do you extract video so you can watch it on other devices?
MATTHEW GAST

73 MONITORING VIRTUAL MEMORY WITH VMSTAT
Need to monitor performance on your systems? vmstat helps you understand when memory is the bottleneck.
BRIAN K. TANAKA

75 MAKING LINUX ACCESSIBLE FOR THE VISUALLY IMPAIRED WITH SPEAKUP
A great article from an amazing 16-year-old kid on a project that's changing his life.
AMEER ARMALY

78 UNIX: OLD SCHOOL
Want to walk down "Virtual Memory" lane? SIMH allows you to emulate historical UNIX implementations, like UNIX V5 from back in 1974.
MATTHEW HOSKINS

82 WIRELESS PORTALS WITH WIFIDOG
Wifidog teaches Wi-Fi hotspots new tricks. Leave your Wi-Fi access point open, but control and monitor its use.
MICHAEL LENCZNER

85 VIM FOR C PROGRAMMERS
Learn to use Vim like the power tool that it is.
GIRISH VENKATACHALAM

89 MINI KDE FOR A LIGHTWEIGHT DESKTOP
See how the RULE Project helps schools and charities run current KDE apps on older hardware using a customized, bare-bones KDE implementation.
MARCO FIORETTI

EMBEDDED

36 A MEMORY MANAGEMENT APPROACH FOR SWAPLESS EMBEDDED SYSTEMS
Avoid the dreaded Out-of-Memory killer exception using these strategies for memory allocation.
MAURICIO LIN, VILLE MEDEIROS, RAONI NOVELLINO, ILIAS BIRIS AND EDJARD MOTA

TOOLBOX

14 AT THE FORGE
Working with ActiveRecord
REUVEN M. LERNER

22 COOKING WITH LINUX
amaroKing the Night Away
MARCEL GAGNÉ

28 WORK THE SHELL
Get Started with Redirection
DAVE TAYLOR

30 PARANOID PENGUIN
Single Sign-On and the Corporate Directory, Part I
TI LEGGETT

COLUMNS

44 LINUX FOR SUITS
The World Live Web
DOC SEARLS

96 EOF
Open-Source Use Accelerates Software Development
PALLE PEDERSEN

DEPARTMENTS

4 FROM THE PUBLISHER
6 LETTERS
10 UPFRONT
68 NEW PRODUCTS
81 ADVERTISERS INDEX
95 MARKETPLACE

Install MythWeb to get a Web front end to view your program guide, scheduled recordings and already recorded programs (page 64).

NEXT MONTH

HOME PROJECTS

Linux has grown up in video production. A professional broadcast technician shows you how easy it is to produce your own video DVD using Kino.

Once you have a few videos under your belt, you will need some serious backup space. Duncan Napier will show you how to build a low-cost, terabyte-sized backup server.

You may have used Skype to make phone calls over your computer. We will show you how to set up a Skype server so you can use your regular phones.

COVER PHOTO BY JOSHUA BLAKE

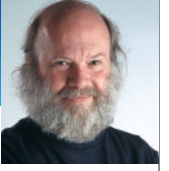

The Desktops Are Coming

Although *LJ* readers have been using Linux desktops for years, putting desktops in front of ordinary users is now a reality. To help, we created *TUX*. **BY PHIL HUGHES**

I know we are all going to miss Don Marti. Many of us *LJ* folks have known Don "forever", and we have worked with him for five years. Don had the big picture, knew his bits and was a great writer—exactly the right mix.

Getting to write this piece gives me a chance to talk a bit about what is changing in Linux and, because of that, what I have been working on.

We as a company decided to switch everyone in the office over to KDE, back on version 1. Everyone thought I was crazy, and many times, I believed them. However, this meant they all knew what Linux was and used it every day. Since then, while most people were watching Linux (and Apache) take over the server market, the desktop quietly matured. It isn't perfect today, but it is certainly easy for your grandmother to sit down at a Linux box and use it. But *LJ*'s job isn't done. Desktops still require new drivers, new applications, security and, in general, administration. *LJ* is here to help you with that, and will be for years, but Linux has a growing *user* market—users like the receptionist I subjected to KDE 1 so many years ago.

More than a year ago, we started working on a new magazine named *TUX*, and it's different in a lot of ways—not just in audience: 1) it is distributed as a PDF; 2) it tells you how to get things done rather than what is inside; 3) it's free; and 4) all the back issues are available for free too.

Is there a catch? Yes. We want lots more people to use Linux. Some of them will become geeks and, thus, *LJ* readers. But, lots of them will simply get to see why we are so excited about what we do and, hopefully, buy a few Linux systems. Some of those people will buy a system for home, but many will end up using Linux at work. That gets us all closer to the goal—World Domination. If you have a friend or relative who just wants to use a computer and you think Linux is the right answer, point them at *TUX* (www.tuxmagazine.com) for articles and free subscription links. And, maybe if you are pretty geeky and know how to do everything on the command line, you should get a subscription too. Although I am writing this with vi, some GUI programs out there are useful—from amaroK to Inkscape.

Enough about what else we are up to. Let's talk about what we did this month in *LJ*.

Reuven continues looking at pieces of Ruby on Rails, focusing on ActiveRecord, the object-relational mapper (page 14). I have been working on a project using Ruby recently, although we rejected using Rails because the project was far from a pure Web application. Ruby on Rails certainly has its place, and Reuven is doing a great job of showing us how to use it.

Marcel looks into amaroK and new features that have recently appeared in this fancy music player (page 22). Even though OGG isn't French for anything, Marcel fills you in on what amaroK can do.

Beyond that, we show you how to make Schenker graphs, master DVDs, replace your TiVo with your own Linux box, squeeze parts of KDE into a small footprint and a whole lot more.

Phil Hughes is Group Publisher for SSC Publishing, Ltd.

LINUX JOURNAL

DECEMBER 2005
ISSUE 140

EDITOR Kevin Bedell, ljeditor@ssc.com

EXECUTIVE EDITOR Jill Franklin, jill@ssc.com

SENIOR EDITOR Doc Searls, doc@ssc.com

WEB EDITOR Heather Mead, heather@ssc.com

ART DIRECTOR Garrick Antikajian, garrick@ssc.com

EDITOR EMERITUS Don Marti, dmarti@ssc.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

PRODUCTS EDITOR James Turner, newproducts@ssc.com

CONTRIBUTING EDITORS

David A. Bandel • Greg Kroah-Hartman • Ibrahim Haddad • Robert Love • Zack Brown • Dave Phillips • Marco Fioretti • Ludovic Marcotte • Paul Barry • Paul McKenney

PROOFREADER Geri Gale

VP OF SALES AND MARKETING Carlie Fairchild, carlie@ssc.com

MARKETING MANAGER Rebecca Cassity, rebecca@ssc.com

INTERNATIONAL MARKET ANALYST James Gray, jgray@ssc.com

REGIONAL ADVERTISING SALES

NORTHERN USA: Joseph Krack, +1 866-423-7722 (toll-free)

EASTERN USA: Martin Seto, +1 905-947-8846

SOUTHERN USA: Laura Whiteman, +1 206-782-7733 x119

INTERNATIONAL: Annie Tiemann, +1 866-965-6646 (toll-free)

ADVERTISING INQUIRIES ads@ssc.com

PUBLISHER Phil Hughes, phil@ssc.com

ACCOUNTANT Candy Beauchamp, acct@ssc.com

***LINUX JOURNAL* IS PUBLISHED BY, AND IS A REGISTERED TRADE NAME OF, SSC PUBLISHING, LTD.**

PO Box 55549, Seattle, WA 98155-0549 USA • linux@ssc.com

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, CEO, Penguin Computing

Bruce Perens

Bdale Garbee, Linux CTO, HP

Danese Cooper, Open Source Diva, Intel Corporation

SUBSCRIPTIONS

E-MAIL: subs@ssc.com • URL: www.linuxjournal.com

PHONE: +1 206-297-7514 • FAX: +1 206-297-7515

TOLL-FREE: 1-888-66-LINUX • MAIL: PO Box 55549, Seattle, WA 98155-0549 USA • Please allow 4–6 weeks for processing address changes and orders • PRINTED IN USA

USPS *LINUX JOURNAL* (ISSN 1075-3583) is published monthly by SSC Publishing, Ltd., 2825 NW Market Street #208, Seattle, WA 98107. Periodicals postage paid at Seattle, Washington and at additional mailing offices. Cover price is $5 US. Subscription rate is $25/year in the United States, $32 in Canada and Mexico, $62 elsewhere. POSTMASTER: Please send address changes to *Linux Journal*, PO Box 55549, Seattle, WA 98155-0549. Subscriptions start with the next issue. Back issues, if available, may be ordered from the Linux Journal Store: store.linuxjournal.com.

LINUX is a registered trademark of Linus Torvalds.

TYAN

Fast, Flexible, and Feature-Rich!

PCI Express and EM64T Servers Have Arrived

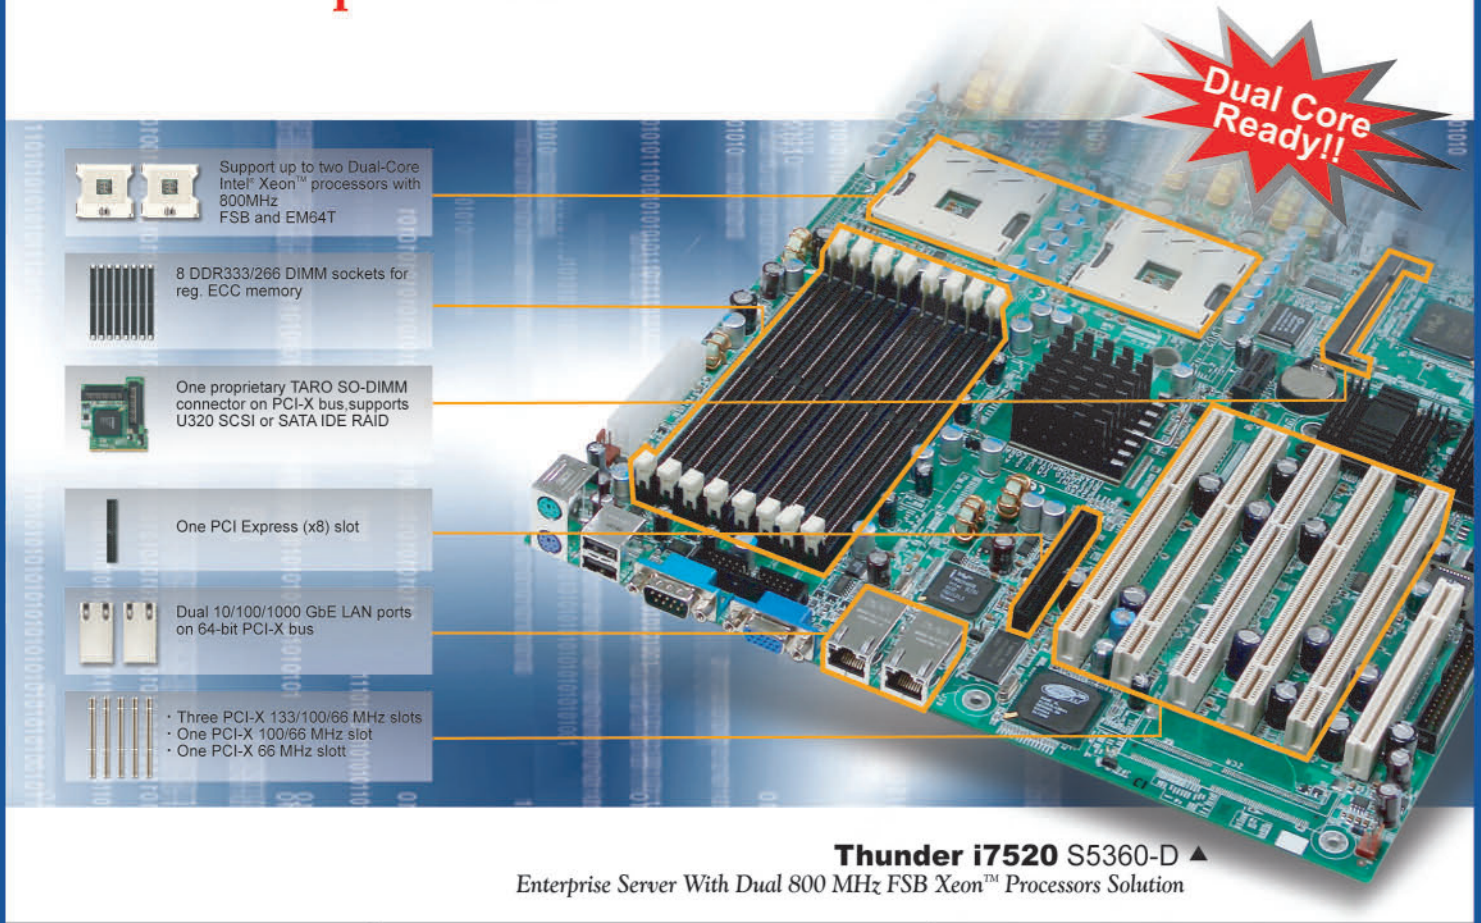

Thunder i7520 S5360-D ▲

Enterprise Server With Dual 800 MHz FSB Xeon™ Processors Solution

intel inside XEON

Thunder i7520 S5360-D

- Supports two Dual-Core Intel® Xeon™ processors with 800 MHz FSB and EM64T
- 8 DIMMs for DDR266/333 memory
- Three PCI-X 133/100/66 MHz slots, one PCI-X 100/66 MHz slot, one PCI-X 66 MHz slot and one 33 MHz PCI slot
- One PCI Express™ x8 slot
- One proprietary SO-DIMM connector on PCI-X bus, supports U320 SCSI or SATA
- Dual GbE LAN

Tiger i7320 S5350-D

- Supports two Dual-Core Intel® Xeon™ processors with 800 MHz FSB and EM64T
- 8 DIMMs for DDR266/333 memory
- Two PCI-X 64/66 MHz slots; three 32/33 PCI 2.3 slots
- One proprietary SO-DIMM connector on PCI-X bus, supports U320 SCSI or SATA
- Dual PCI Express GbE LAN

TYAN COMPUTER CORP.

Tyan Computer USA
3288 Laurelview Court
Fremont, CA 94538 USA
Tel: +1-510-651-8868 Fax: +1-510-651-7688
Pre-Sales Tel: +1-510-651-8868 x5120
Email: marketing@tyan.com

For more information about this and other Tyan products, please contact Tyan Pre-Sales at (510) 651-8868 x5120, or contact your local Tyan system integrator/reseller.

www.tyan.com

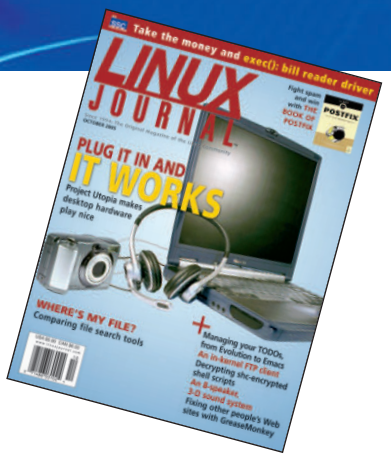

Questions on Internet Radio to Podcast Article

Great article in *LJ*, September 2005 [“Internet Radio to Podcast with Shell Tools” by Phil Salkie]! There is a slight problem: in your final bash script you sleep for 2.1 hours and leave for the script to then figure out what the process id number of mplayer is. More than two hours seems to be quite an extended period of time on which any process may be forked onto the background, and then that subsequent process would be the one killed. Why don’t you capture the PID of mplayer onto a local variable and then kill that PID stored on that variable?

--

Juan C. Müller

Error in Kernel Korner

There is an error in the code samples for the Kernel Korner article [“Sleeping in the Kernel” by Kedar Sovani] in the September 2005 issue: wait_event() and wait_event_interruptible() should not be passed the address of my_event, but my_event itself. That is because they are macros, and their implementations will wind up using the address-of operator (&) to take the address of the parameter they are passed.

--

Bob Bell

On Patents

In Don Marti’s editorial reply to Darin Riedlinger’s letter “Multimedia Lock-in?” he states, “You can create your own media in patent-free formats you can use on any OS.”

I seriously doubt this. If, for example, OGG would become very popular such that MP3 players (of the hardware version) would start to come out without MP3 support, with the intent of not paying the royalties for the MP3 patents, then those holding the patents on MP3 would be quick to find a patent that also applies to OGG.

Currently, patents are granted on way too obvious things. “Audio compression by omitting nuances that the human ear cannot detect” is a description I heard of a patent that Fraunhofer supposedly holds. On an authoritative-looking Web site I found titles as short as: “method for coding an audio signal”, which with a bit of fantasy can really apply to, say, OGG.

Thus, a statement that implies that media formats like OGG are patent-unencumbered cannot be made. The only thing that you can say is that nobody has stepped forward to claim that he or she owns a patent on something in a format like OGG.

The intent of patents has always been to protect:

1. The small inventor who invents something that nobody would have invented, but is somewhat obvious after the fact (for example, a chain with differently sized gears to drive the rear wheel of a bicycle). An inventor like this may need some time to set up a factory and earn a fair compensation for his “brilliance”.

2. The big companies who spend big bucks to develop something interesting. These need a “grace period” to earn back their investment.

The whole patent application process has become too expensive for the first type of inventors. And the big corporations are claiming that they actually *do* have millions of “inventions” that warrant the second type of protection. But way too many “the time is right” type of things are being patented.

I’m convinced that any serious application, open source or not, will violate several patents. If some smaller company happens to have patented something later seen in say Microsoft Word, then they might get up the nerve to step up to Microsoft and ask for royalties. In return, Microsoft will research whatever the smaller company is making and try to find a patent infringement on something they hold. Most likely it will turn into a “we won’t pay for the use of your patent in return for the use of ours.”

Big companies have an arsenal of patents they can use for this type of stuff. Remember when IBM first was approached by SCO? Within a month, IBM had found a bunch of patents of theirs that SCO was violating. Only if the other party becomes “annoying” do the big patent monsters come out of hiding and start to throw threats around.

Becoming “annoying” can be done in several ways: cutting into royalty payments on another patent (MP3/OGG), asking for royalty payments for some obscure patent (small company/Microsoft example) or filing a big lawsuit (SCO/IBM).

--

Roger Wolff

ALSA Problem

In the article “A User’s Guide to ALSA” in the August 2005 issue, Dave Phillips mentioned having a desktop system with a SoundBlaster Live! Value sound card. This caught my attention because I have the same kind of sound card in my system. I have been unable to use ALSA, however, because I have digital speakers and have been unable to determine how to tell ALSA to switch my card to digital output. I am able to switch to digital output under OSS using a utility from the emu10k1 package available at sourceforge.net/projects/emu10k1. The actual command line that I use is `emu-config -d` with the -d meaning “switch output to digital”. I would like to begin using ALSA, however, because it appears that development on the emu10k1 package has been discontinued, and the days for OSS appear to be numbered. Perhaps Mr Phillips or one of your readers might have an answer to my dilemma.

--

Mark Iszler

Dave Phillips replies: *the SBLive is certainly a complicated beast. Alas, I don’t have digital speakers, so I can’t provide a direct answer to your question. However, I suggest checking your mixer for channels named IEC958-whatever. These are the*

Reclaim lost time

Levanta Intrepid ™

The world's first Linux management appliance

Plug the Levanta Intrepid ™ into your network and perform the most important Linux management tasks in a fraction of the time you spend now. And gain power and flexibility that you've never had before:

- **Fast & Portable:** Provision servers or workstations practically anywhere, anytime – in minutes. Swap them around, mix it up.
- **Flexible:** Supports commodity hardware, blades, virtual machines, and even mainframes.
- **Out of the Box:** Includes pre-defined templates for servers, workstations, & software stacks. Or create your own.
- **Total Control:** Track any file changes, by any means, at any time. And undo them at will.
- **Disaster Recovery:** Bring dead machines quickly back to life, even if they're unbootable.

Based upon technology that's already been proven in Fortune 500 enterprise data centers. Now available in a box, priced for smaller environments. **Just plug it in and go.**

30-Day
Money-Back Guarantee
Order online by 12/31/05
Get $500 Off
Enter PROMO CODE: LJ1205

LEVANTA®
www.levanta.com
1.877.LEVANTA

WINNER
LinuxWorld CONFERENCE & EXPO
Most Innovative Hardware Solution

© 2005 Levanta, Inc. All rights reserved. Levanta and the Levanta logo are registered marks of Levanta, Inc.

SBLive digital (S/PDIF) channel controls as they appear in alsamixer and in qamix.

Be sure your digital speakers are connected to the digital output of the card. (Sorry, just trying to be complete.)

I'm a little unclear as to whether you're actually using ALSA yet. Also, let me know what kernel version and ALSA release you're using; this makes a difference.

More on Linux Hardware Support

Robert Love's article "Project Utopia" [October 2005] is a great overview of the direction in which Linux hardware support is moving.

We have developed some programmable network hardware for the PCI bus, and are in the process of developing Linux drivers for it. We would love to see more articles on HAL, udev and ssyfs, explaining how they fit together, examples of how to use them and so forth.

--
Greg Watson

And More Requests

I would love to see more articles about installing and tuning a Linux Debian distro on a PPC box. I dual-boot into OS X 10.3.9 and Ubuntu on a G3 iBook and have yet to find anyone on-line who knows how to get ALSA drivers working on it.

I love the Linux platform and realize the imminent move to the Intel hardware might make some of this moot, but there are people out there right now with PPC machines who want to do sound and video using free Linux tools. So, how about an article or column dedicated to getting ALSA drivers working on a PPC? Keep up the great work—love *LJ*!

--
Kim Cascone

Getting Organized

I appreciated the article in the October 2005 *Linux Journal* by Sacha Chua ["Taming the TODO"]. Emacs has been on my list a long time, but I still haven't started using it; I just got started with vi and didn't want to learn something new. Maybe I'll give Emacs a try now.

I also appreciated your acknowledgement of the index card method. I use that a lot, especially when at a customer site without my laptop or guaranteed Internet access.

Another solution you didn't mention is the wiki. Although not as formal or organized as an issue-tracking system, it does revision control on your documents and makes them publicly accessible via a Web browser. I use wiki for my personal TODO list as well as communal TODO lists for several projects.

--
SamU

Sweet!

Here is a photo of a bag of sweets, popular here in Thailand. Comes in menthol flavour as well!

Great magazine by the way.

--
Andrew

Open the Name Linux

One indication that Linux(R) will have reached mainstream would be if there were so many companies supporting Linux listed in the Yellow Pages that the phone company has to create a separate Linux category. However, this ideal seems distant, because attractive Linux company names are being declined by the Linux Mark Institute (LMI).

One recent example of an unacceptable name was discussed on our local Linux mailing list. An entrepreneur learned from LMI that "Linux of Sacramento" was unacceptable—a name likely to generate many phone calls.

An LMI representative said that they are assigned the responsibility to protect the health of the Linux mark by keeping it from being diluted. When I asked the representative for an explanation of how a mark could become unhealthy by dilution, his explanation was too obtuse for me to understand. However, I could understand that a Linux name license would be approved if the name did not imply an exclusive source of Linux in an area. But how would an exclusive source of Linux make the name unhealthy?

I pressed the representative for a case study of a trademark becoming unhealthy by dilution—Kleenex(R) or Xerox(R)? No, they may become generic—another issue entirely.

Linux will become mainstream when the Linux mark saturates the public. Efforts to prevent saturation is counterproductive; instead, decision-makers should consider further opening up the Linux name.

--
Tim Riley

Errata

Regarding the article "The Ultimate Linux Lunchbox" by Ron Minnich in the November 2005 issue of *LJ*: the first mini-cluster by Sandia was the brainchild of both Rob Armstrong and Mitch Williams. This system was built in 2001, not 2000 as was stated in the article.

--
Ron Minnich

We welcome your letters. Please submit "Letters to the Editor" to ljeditor@ssc.com or SSC/Editorial, PO Box 55549, Seattle, WA 98155-0549 USA.

MBX is the Industry Leader for Server Appliances

From design to delivery we are dedicated to building a partnership with you.

MBX RP-1013 Platform

- Intel® Pentium 4 630 Processor at 3.0 GHz
- 1U Rackmount Chassis 16.5" Deep
- 1GB PC4200 DDR2 Memory
- Maxtor 80GB Serial ATA Hard Drive
- Eight Gigabit NIC's, 4 ports with by-pass
- Optional 16x2 LCD with Keypad
- On-board Compact Flash Socket
- Branded With Your Color and Logo
- Custom OS and Software Install
- No Minimum Quantity Required
- 3 Year Warranty

$1,699 or lease for **$52**/mo.

MBX RP-2012 Platform

- Intel® Xeon Processor at 3.0D GHz
- 2U Rackmount Chassis
- 1GB PC3200 DDR Memory
- Maxtor 73GB SCSI Hard Drive
- Dual On-board Gigabit NIC's
- Custom OS and Software Install
- No Minimum Quantity Required
- 3 Year Warranty

$2,899 or lease for **$89**/mo.

MBX is the leader in custom appliances. Many premier application developers have chosen MBX as their manufacturing partner because of our experience, flexibility and accessibility. Visit our website or better yet, give us a call. Our phones are personally answered by experts ready to serve you.

MBX systems

www.mbx.com

1.800.939.0971

Intel, Intel Inside, Pentium and Xeon are trademarks and registered trademark of Intel Corporation or its subsidiaries in the United States and other countries. Lease calculated for 36 months, to approved business customers. Prices and specifications subject to change without notice. Setup fee may apply to certain branding options. Motherboard Express Company. 1101 Brown Street Wauconda, IL. 60084.

On the WEB

We at LinuxJournal.com have been fortunate over the years to receive all sorts of how-to and DIY articles. Our authors love to write about their cool projects—you know, the stuff you guys are piecing together in basements and workshops with soldering guns, breadboards, microcontrollers and a few lines of C written in vi. And LinuxJournal.com readers love to read about what other people are doing so they can hack a project for their own needs. Well, we want more of this exchange. So, we're asking our readers to tell us what they're building, hacking and conjuring. Send your project outlines and article proposals to webeditor@ssc.com.

In the meantime, LinuxJournal.com offers these project articles to hold you over:

- "Learning to Master MythTV" (www.linuxjournal.com/article/8564) by Colin McGregor starts by explaining what MythTV is and when it's a good idea to build your own personal video recorder and then moves on to explore MythTV plugins for DVDs, photo galleries, games and more.

- Although not a basement project, the FreeNX Project is both cool and useful. Kurt Pfeifle, a member of the FreeNX Development Team, offers a seven-part series that introduces NX technology and explains how it lets you run remote X11 sessions across slow or low-bandwidth network connections. In Part 5 (www.linuxjournal.com/article/8538), Kurt provides step-by-step instructions for maneuvering your way though the NX interfaces.

We recently posted the 2006 Editorial Calendar on LinuxJournal.com; it's available at www.linuxjournal.com/xstatic/author/topicsdue. It lists the focus topic we have planned for each issue in 2006. Take a look at the topics—ranging from "Home Projects" to "Building Dynamic Web Sites"—and send a proposal to ljeditor@ssc.com if you have an idea for an article.

diff -u

What's New in Kernel Development

SMBFS has been orphaned. **Urban Widmark**, the official maintainer, has stopped responding to e-mail about the filesystem, and **Adrian Bunk** has put out the call for someone to step up and maintain this code. The situation is colored by the fact that **CIFS**, a potential replacement, does not yet support the full array of Windows variants covered by SMBFS. Apparently **Red Hat** discovered this when they tried to remove SMBFS in Fedora and had to re-enable it fairly quickly. With the CIFS developers working to extend the number of supported systems, the situation of SMBFS is even more uncertain. Should a new maintainer come forward? Should the code just sit quietly until it can be replaced by CIFS? The future of this corner of the kernel seems yet to be decided.

The **linux-kernel mailing list** has received an infusion of life. **Dell** recently donated a powerful computer to host the list, and the result has been much better latency between the time a user posts to the list, and the time readers receive that post. Over the years, as the number of silent readers and active posters has gone up and up and up, the hardware running linux-kernel (and the rest of the vger mailing lists) has occasionally been overwhelmed. Various companies always have offered generous donations when speed or bandwidth has gotten tight to keep these lists running properly. Dell's gift, and Red Hat's donation of a 1 gigabit network connection, should ensure linux-kernel's smooth operation for the near-to-mid future.

Michael S. Tsirkin has gone through the kernel sources, identifying and documenting the basic stylistic standards for **whitespace usage**. He started this project as a way to help his coworkers get started with kernel development, but published the results when he realized they might actually have a wider appeal. A set of kernel coding standards already exists in the Documentation/CodingStyle file distributed with the official sources, but that file neglects to cover much of the intricate details of whitespace usage. Michael's document is a first. As soon as he posted it, a bunch of other developers offered detailed suggestions and refinements, so the latest version is probably quite reliable.

Andrea Arcangeli has written a tool to help track how many people actually test each new kernel. This tool, called **klive**, runs in user space on the computers of willing participants and reports various system statistics to Andrea's server at klive.cpushare.com, where the results are aggregated and displayed. So far, more than 100 users are participating in the effort. One problem various kernel developers have with this project is the possibility that users might think of it as a tool to spy on them. As a result, it is less likely that Andrea will be able to migrate his tool to a full-kernel feature. Probably, klive will remain just a user program, unless developers' concerns can be clearly assuaged.

Adrian Bunk, always on the lookout for ways to clear out kernel deadwood, has been pushing a patch to remove support for older GCC versions. According to Adrian, newer compilers are perfectly able to compile the kernel, and continuing to support the older compilers results in a lot of conditional code that makes the kernel uglier, larger and harder to maintain in some areas. Nevertheless, it seems that many kernel developers feel quite strongly that at the very least, **GCC 2.95** must continue to be supported. GCC 2.95 is blazingly fast compared to recent compilers, and anyone compiling multiple kernels per day (as kernel developers are wont to do) saves considerable time by relying on GCC 2.95 instead of the more recent compilers. So it looks as though Adrian's patch may have to wait until newer compilers can better compete for speed.

Chris Wedgwood recommends boycotting **NVIDIA** until they start releasing the specifications needed to write open-source drivers for their hardware. This came up recently when **Michael Thonke** asked whether Linux would implement **NCQ support** for NVIDIA NForce4 (CK804) SATAII-based chipsets. **Jeff Garzik**'s reply was that there were no plans to implement this because there was no documentation from NVIDIA. He also said, "They are the only company that gives me zero information on their SATA controllers." With NVIDIA apparently so hostile to free software, Chris argues, it's up to the rest of us to send them a message by not purchasing their hardware until they change their tune.

—ZACK BROWN

The Power of Choice...

Command the game with your next I/O move.

Modularity. Scalability. Reliability. Cost-effectiveness. These represent the solid foundations that SBE delivers to OEMs for building innovative end solutions. Partnering with SBE for networking and communications I/O solutions allows you to take advantage of proven technology and field-tested products designed to optimize performance for your unique application needs.

SBE offers a full spectrum of interface cards, ranging from T1 and T3 to Gigabit Ethernet and IPsec/SSL acceleration. These boards are available in multiple form factors, including PCI, PMC, and PTMC. Customers have the choice of buying these boards individually or bundling any of the PMC/PTMC modules with our intelligent core processing platforms to create a flexible, cost-efficient blade solution ideal for serving demanding telecom applications. Full Linux support is available on every board.

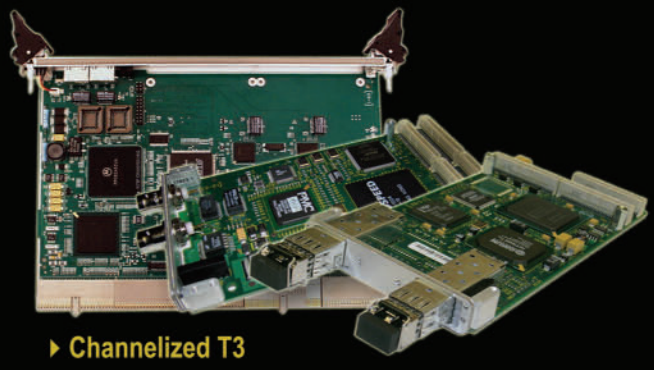

- Channelized T3
- 24-port T1/E1
- LAN/Ethernet
- Storage
- IPsec/SSL Encryption
- Blade platforms
- I/O and beyond...

sbe®

Linux *On Demand*

flexibility on demand | 925-355-2000 | info@sbei.com | www.sbei.com

They Said It

At Parrs Wood OSS is seen not as merely a way of saving money, but rather of spending it more effectively.

—BBC NEWS: news.bbc.co.uk/1/hi/education/4642461.stm

When I switched from Windows to GNU/Linux (Red Hat/Fedora/Debian mostly) about five years ago, I found a vast developer's playground. It was like the old days of Compuserve, which was a candy aisle of freeware. Free software is still like that for me; there's a lot of it to explore, and I can see the source code without significant restriction. I can *use* the source, and I can share the source...which is something geeks love to do. The Windows world by the mid-1990s was very closed (still is mostly), something that's really restrictive as a developer.

—ANONYMOUS, ON IT GARAGE: www.itgarage.com/?q=node/617#comment

You have reached the pinnacle of success as soon as you become uninterested in money, compliments or publicity.

—THOMAS WOLFE, *THE SUN*, JULY 2005

There is no such thing as a "personal" blog if you are employed.

—CHRIS DIBONA

Money can't buy happiness, but it can buy a Linux box.

—JON WATSON, www.jonwatson.ca/blog

If AOL ruled the world, they would slap training wheels on skateboards and charge kids $20/month to go slower and to be able to do fewer things.

—TONY PIERCE, www.tonypierce.com/blog/2003_07_13_blogarc.htm

Today's laptops have become obese. Two-thirds of their software is used to

100 Million × $100 Linux Laptops

MIT's Media Lab is developing a $100 US Linux-based laptop that will "be able to do most everything except store huge amounts of data". The units will have color displays, Wi-Fi, mesh networking, cell-phone connectivity and "USB ports galore".

Nicholas Negroponte, chairman and co-founder of the Media Lab, announced the initiative in January 2005 at the World Economic Forum in Davos, Switzerland. Details of the initiative were published in August 2005.

In a Q&A that ran with the August 2005 announcement, Negroponte said, "...we will market the laptops in very large numbers (millions), directly to ministries of education, which can distribute them like textbooks." He also calls the project "One Laptop Per Child". The plan is to have units ready for shipment by late 2006 or early 2007. The goal is to produce and distribute 100 million of them.

Tom Limoncelli, co-author of *The Practice of System and Network Administration*, said, "The thought of laptops distributed like textbooks could be as revolutionary for spreading hardware as Linux was for spreading UNIX-like systems."

See laptop.media.mit.edu.

—DOC SEARLS

PROPRIETARY DEVICE DRIVERS?

If you've been reading *Linux Journal* for a while, you'll notice that everyone here tells you to stay away from proprietary device drivers. Video cards, wireless network hardware and Fibre Channel hardware have been especially problematic.

By releasing a proprietary driver, not only does a vendor shut itself out of the non-x86 embedded market and pass up free driver testing and optimization from the experts on the linux-kernel mailing list, it's also hurting itself with regular Linux customers too.

Here's what readers said in a survey (numbers rounded):

- **We don't use proprietary drivers on Linux: 20**
- **We'll use a proprietary driver only if there's no competing hardware with a GPL driver: 14**
- **A proprietary driver tends to make us less likely to buy a piece of hardware, but doesn't rule it out: 35**
- **We'll use proprietary drivers only if our Linux hardware vendor or distribution vendor commits to supporting them: 8**
- **Whether the driver is GPL or proprietary doesn't matter in our hardware buying decisions: 20**
- **We prefer proprietary drivers to GPL drivers: 0**

That last one is there for the marketing guy at an "enterprise" hardware vendor who told me that the company's enterprise customers would never want GPL drivers for their GPL OS. Sounds like you need to get out and talk to the customers a little more, dude.

One support engineer at a popular enterprise distribution told me that his group has to support some proprietary drivers, but that when those drivers lead to support calls, the customers ask about alternative hardware with GPL drivers. With the Linux hardware market at more than $4 billion a year, letting your lawyers slap a restrictive license on your drivers could be an expensive mistake.

—DON MARTI

routerOS: routerBOARD.com

User Management

- support more than 3000 PPPoE or HotSpot clients
  - full RADIUS support for user parameters
    - tx/rx speed, address, filter rules
    - supports RADIUS real time modification of parameters while users are online
- Peer to Peer protocol control (P2P)
  - per client P2P tx/rx rules with burst support
  - P2P pool
  - complete blocking of P2P

Wireless AP and Backbone

- Wireless monitoring
  - Frequency scanning with detailed report
  - Raw wireless packet sniffer
    - streaming option to Ethereal analyzer
    - option to save to a file format supported by Ethereal
  - Snooper packet inspection
    - analyzes all raw frames received for wireless parameters
    - monitors a single channel or all channels

- Nstreme wireless polling protocol
  - no decrease in speed over long distances (as seen with the 802.11 ACK packet bottleneck)
  - polling improves speed and eliminates contention for access to the wireless bandwidth
  - access point control over Nstreme clients tx data to optimize use of the wireless medium
  - RADIUS support for the access control list including bandwidth settings for wireless clients

- Full 802.11a/b/g support

The above is a brief description of a few features,

RouterBOARD 500 $140

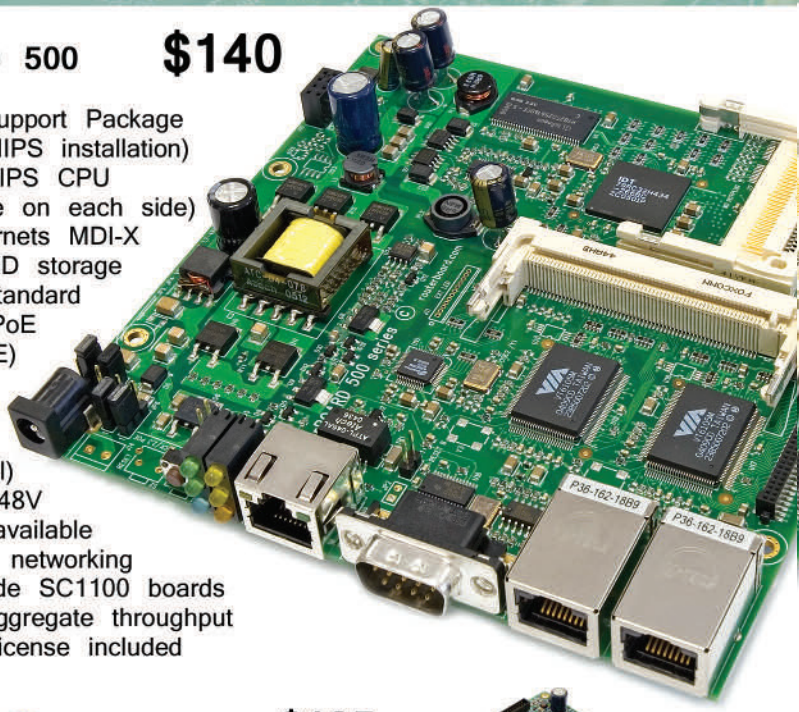

- Linux Board Support Package (full Debian MIPS installation)
- 266-400MHz MIPS CPU
- 2 miniPCI (one on each side)
- 3 10/100 Ethernets MDI-X
- 64/128MB NAND storage
- PoE 802.3af standard and passive PoE (also 12V PoE)
- Compact Flash
- Low power
- 32MB DDR (64MB optional)
- 6-24V and 25-48V power mode available
- 2-3x faster for networking than the Geode SC1100 boards
- 200-300Mb/s aggregate throughput
- L3 RouterOS license included

$95

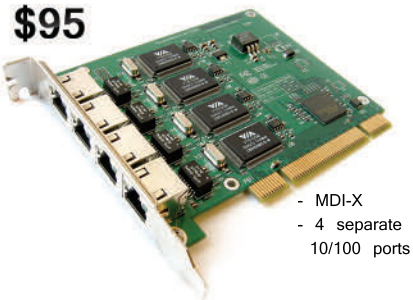

- MDI-X
- 4 separate 10/100 ports

RouterBOARD 44

For the Router Builder !

- rackmount servers and routers
- up to 24 Ethernet ports in a PC
- no more straight/cross cable problems
- server quality VIA VT6105 chips

$195

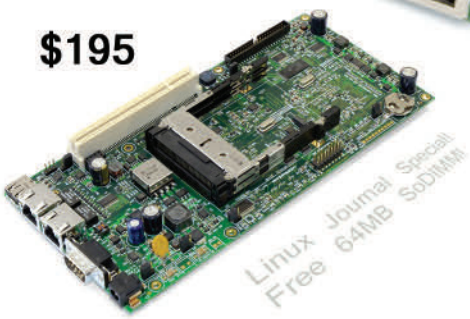

RouterBOARD 230

No feature left behind !

Integrated router with various interfaces. Use as an AP on a tower with up to 500ft PoE. Includes IDE/CF, miniPCI, USB, PCMCIA, UART, PCI, GPIO, LCD controller, Linux SDK, and more.

$120 $65

Eight ports Four ports

RouterBOARD 11/14/18

Multi radio tower !

MiniPCI to PCI adapters for multi radio system. Tested with sixteen radios in one Router/AP.

$240

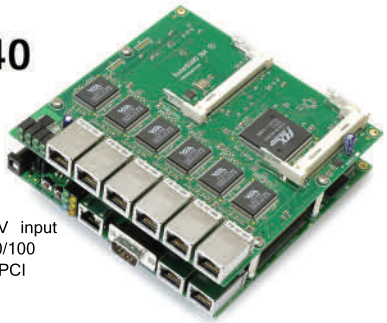

- PoE
- 10-56V input
- 9x 10/100
- 6x mPCI

RouterBOARD 500 & RouterBOARD 564

The Wireless Switchboard !

For a complete multi-radio tower system, the RouterBOARD 500 can carry a daughterboard (RouterBOARD 564) which adds six ethernets and four miniPCI.

MikroTik User Meeting 2006
Prague, Czech Republic, January 19-20

contact sales@routerboard.com or go to www.routerboard.com
all copyrights and logos belong to their respective owners

Working with ActiveRecord

Reuven Lerner continues his series on Ruby on Rails with this look at database integrity checking using ActiveRecord. BY REUVEN M. LERNER

For the past few months, we have been looking at Ruby on Rails, the hot new open-source toolkit for creating Web/database applications. One of the core elements of this toolkit, as we saw last issue, is the ActiveRecord class, which automatically translates between Ruby objects and data in a relational database. Object-relational mappers, as such software is often known, bridges the gap between the object-oriented and relational worlds, which treat data in fundamentally different ways.

This month, we look at some of the ways we can modify ActiveRecord to validate our data in various ways. We also see how we can work with classes that depend on one another, doing something a bit more sophisticated than the basic scaffolding provides with only a few simple lines of code.

Primary Keys

When I first started to work with relational databases, I would create tables that looked like this:

```
CREATE TABLE People (
    first_name    TEXT  NOT NULL,
    last_name     TEXT  NOT NULL,
    phone_number  TEXT  NOT NULL,
    email_address TEXT  NOT NULL
);
```

And of course, the above definition of People will work just fine, providing the basis for a computerized address book. However, the above definition has several problems. To begin with, what happens if there is more than one person with the same name? That is, if we have two people named George Washington in our database, we're going to have a serious problem. How will we know which is the George we want?

The solution to this problem is to assign a unique number to each record in the database. Each relational database product has a different way of accomplishing this. In PostgreSQL, we add a new column and assign it a SERIAL type, indicating that it should be a nonrepeating integer:

```
CREATE TABLE People (
    id            SERIAL   NOT NULL,
    first_name    TEXT     NOT NULL,
    last_name     TEXT     NOT NULL,
    phone_number  TEXT     NOT NULL,
    email_address TEXT     NOT NULL
);
```

We then tell PostgreSQL that it should consider id to be not just another column, but the primary key, an identifier that is guaranteed to be unique and that can serve as identification for one row in the table:

```
CREATE TABLE People (
    id            SERIAL   NOT NULL,
    first_name    TEXT     NOT NULL,
    last_name     TEXT     NOT NULL,
    phone_number  TEXT     NOT NULL,
    email_address TEXT     NOT NULL,

    PRIMARY KEY(id)
);
```

Although we can now find people in our address book with their first or last names, we also can do so using their unique ID. Even if there are 100,000 people named George Washington in our database, we can unambiguously find the one that interests us using the id column. Think of the times you have been asked to identify yourself using a driver's license number, a national ID number or a Social Security number, and you quickly will realize that each of these can be used as a primary key in a database.

One additional result of this constraint is that the database creates an index for the id column. Even if you have a very large table of addresses, the fact that id is indexed means that the database can use it to find records quickly. In addition, although SERIAL columns can be set manually in an INSERT statement, just like INTEGER columns, they're normally not set explicitly at all. Rather, PostgreSQL assigns the next consecutive integer to be the column value—perfect for a primary key, whose value must be unique.

Foreign Keys

Primary keys are useful in this way, but we have not yet begun to understand their power. That's because primary keys really come into their own when they make it possible for us to link tables together. For example, consider a computerized appointment calendar that we might want to build as an add-on module to our existing address book. We could create a table like the following:

```
CREATE TABLE Appointments (
    id         SERIAL    NOT NULL,
    person_id  INTEGER   NOT NULL,
    start_at   TIMESTAMP NOT NULL,
    end_at     TIMESTAMP NOT NULL,
    comment    TEXT,

    PRIMARY KEY(id)
);
```

The above table has an id column, uniquely identifying every appointment. It also has two columns identifying the time at which the appointment starts and ends, as well as room for an optional comment or description.

EmperorLinux

...where Linux & laptops converge

You choose your laptop from a wide selection of top tier laptops manufactured by Dell, IBM, Lenovo, Sharp, and Sony. **You choose your distribution** from among the most popular Linux distributions. We'll install your distribution; add our custom kernel; then configure your laptop for full hardware support, including: X at the native resolution, wireless, power management, 3-D graphics, optical drives...

The Toucan: 5 lb Linux

- ThinkPad T series by Lenovo
- 14.1" SXGA+/15.0" UXGA
- X@1400x1050/1600x1200
- ATI FireGL graphics
- 1.6-2.26 GHz Pentium-M 7xx
- 40-100 GB hard drive
- 512-2048 MB RAM
- CDRW/DVD or DVD-RW
- ACPI suspend/hibernate
- Ask about our Sharp laptops

The Rhino: 7 lb Linux

- Dell Latitude D810/M70
- 15.4" WUXGA screen
- X@1920x1200
- NVidia Quadro or ATI Radeon
- 1.73-2.26 GHz Pentium-M 7xx
- 30-100 GB HD (7200 rpm)
- 256-2048 MB RAM
- CDRW/DVD or DVD±RW
- 802.11a/b/g wireless, GigE
- Ask about our Sony laptops

Let EmperorLinux do the rest. Since 1999, EmperorLinux has provided pre-installed Linux laptops to universities, corporations, and Linux enthusiasts. We specialize in the configuration of Linux on the finest laptop and notebooks. We offer a range of the latest Linux distributions a boot options. All systems come with one year of Linux technical suppor manufacturers' warranty.

Introducing the Raven X41 Tablet

Linux Tablet Features

- Pen/stylus input to screen in X
- Handwrite commands to screen focus
- Handwrite notes and convert to text
- Pressure sensitive stylus in GIMP
- Rotate screen orientation

Linux Laptop Features

- ThinkPad X series by Lenovo
- 12.1" XGA w/ X@1024x768
- 1.5 GHz Pentium-M 758
- 60 GB hard drive
- 512-1536 MB RAM
- 802.11a/b/g wireless, GigE
- ACPI hibernate

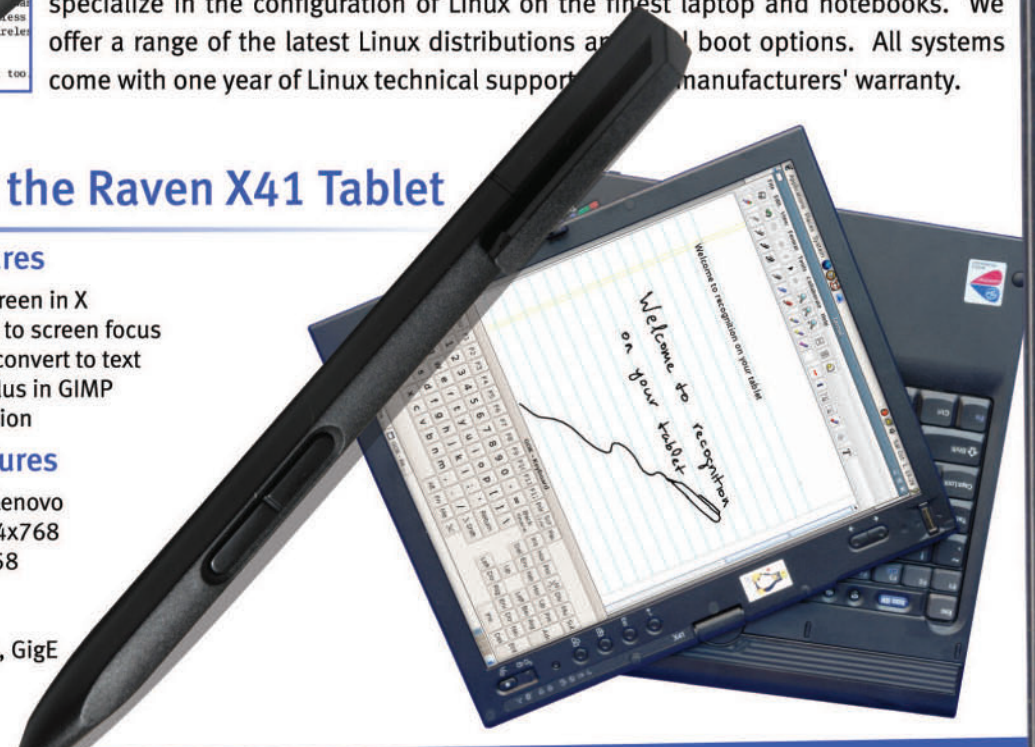

www.EmperorLinux.com **1-888-651-6686**

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

But there is also a person_id column, which allows us to indicate with whom we will be meeting. This database design has a number of problems, but perhaps the most striking one is that there is no constraint (other than NOT NULL) on the value that we can assign to person_id. Even if our People table is empty, we can assign person_id to be 10, 100 or 996—these numbers might be acceptable technically, but they don't help us ensure that person_id refers to an actual person.

The solution is to define person_id as a foreign key, indicating that values of person_id are legitimate only if they reflect an existing value in the People table. In PostgreSQL, we accomplish this as follows:

```
CREATE TABLE Appointments (
    id         SERIAL    NOT NULL,
    person_id INTEGER    NOT NULL  REFERENCES People,
    start_at   TIMESTAMP NOT NULL,
    end_at     TIMESTAMP NOT NULL,
    comment    TEXT,

    PRIMARY KEY(id)
);
```

With these conditions in place, we can be sure that we will be able to make an appointment only with someone in our address book. What happens if we try to get around it? Let's see:

```
INSERT INTO People (first_name, last_name,
                    phone_number, email_address)
    VALUES ('George', 'Washington',
            '202-555-1212', 'first.prez@whitehouse.gov');
```

When we SELECT the elements of our database table, we can see the value that was automatically assigned to our id column:

```
id | first_name | last_name  | phone_number |       email_address

---+------------+------------+--------------+--------------------------
 1 | George     | Washington | 202-555-1212 |  first.prez@whitehouse.gov
```

Now let's insert an appointment with George:

```
INSERT INTO Appointments (person_id, start_at, end_at, comment)
VALUES (1, '2005-Oct-2 18:00', '2005-Oct-2 20:00', 'Dinner');
```

So far, so good. But, what happens if we try to insert an appointment with a nonexistent person?

```
INSERT INTO Appointments (person_id, start_at, end_at, comment)
VALUES (200, '2005-Nov-2 18:00', '2005-Nov-2 20:00',
            'Dinner with no one');
```

PostgreSQL rejects our INSERT statement, saying that inserting the row would violate the constraint introduced with the REFERENCES command:

```
ERROR:  insert or update on table "appointments" violates foreign key
constraint "appointments_person_id_fkey"
DETAIL:  Key (person_id)=(200) is not present in table "addressbook".
```

What happens if we try to remove George from our People table while we have an appointment with him?

```
DELETE FROM People WHERE id = 1;
```

Once again, PostgreSQL rejects our request, indicating this time that we cannot remove an item that is being pointed to:

```
ERROR:  update or delete on "addressbook" violates foreign key
constraint "appointments_person_id_fkey" on "appointments"
DETAIL:  Key (id)=(1) is still referenced from table "appointments".
```

ActiveRecord and Foreign Keys

All of the constraints we have seen so far have been at the level of the database, rather than any application using that database. This potentially means trouble for the users of those applications who don't have access to the database definitions. After all, what is supposed to happen if the application tries to insert, delete or modify a row such that it violates a constraint?

The simple answer, and one that is still prevalent in a surprisingly large number of Web/database operations, is that the program simply reports an error. (Sometimes it even will indicate what the error was, needlessly exposing the offending SQL statement for everyone to see.) In some cases, the application indicates that there was a database problem, or something of the sort.

But, what we really would like is to avoid those sorts of database problems altogether. We would prefer to have the constraints in our database somehow be propagated to the application level, letting the application catch problems before they ever get to the database level.

Although ActiveRecord cannot do this, it comes very close, making it almost trivially simple for us to represent relationships between tables in a Rails application. Let's now create a simple Rails application that uses ActiveRecord to keep track of our address book and calendar information.

We begin by creating the skeleton Rails application by typing `rails addressbook`, which creates an addressbook directory and puts everything underneath that. Then, we modify config/database.yml to point to development, testing and production databases in the appropriate place. (See last month's At the Forge for an example of what database.yml should look like.)

Now, let's create basic models, controllers and views for the People and Appointment tables. We could use the script/generate program that comes with Rails to create them separately. But in many cases, it's easiest to create a barebones application, or scaffold:

```
ruby script/generate scaffold Person
ruby script/generate scaffold Appointment
```

We can now start the test server on port 3000 (script/server); going to /People shows the current list of people and lets us create a new person. Click on the new person link, and you will see the page the scaffolding created. However, not all is perfect here—what happens if you click on the create button at the bottom of the page without entering anything in the text fields?

Assuming the definition of the People table described earlier, Rails will create a new person whose fields are all the

WARNING:

TO THOSE CONSIDERING PERFORMANCE ENHANCEMENT

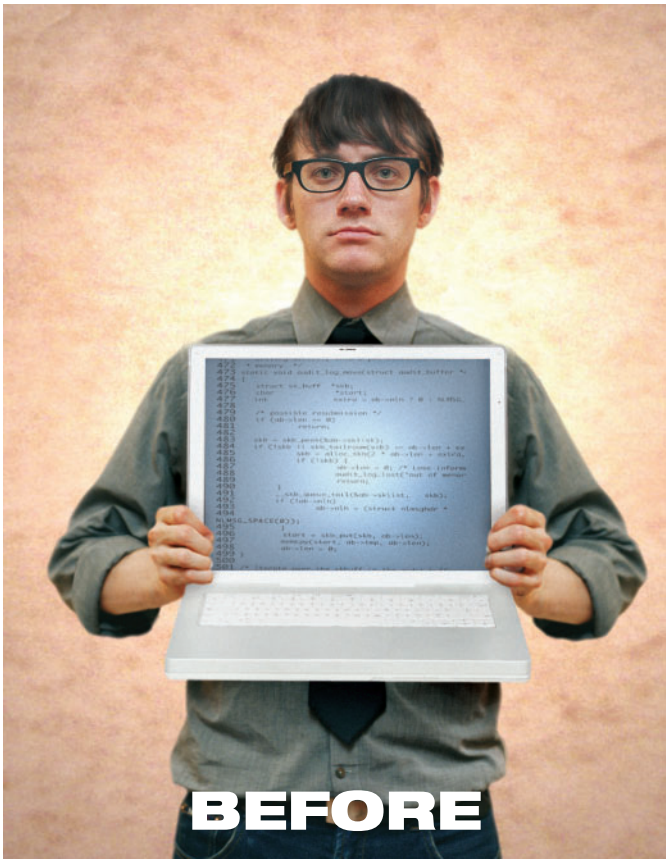

Natural ability and ordinary debuggers can take you just so far. That's why you need NightStar LX™. An integrated suite of tools that gives you full visibility into your Linux® application. You can debug, monitor, analyze and tune at application speed, so you see real execution behavior. Plus, you'll reduce test time and lower costs. NightStar LX. Experience real power for a change.

The NightStar LX suite is extremely habit-forming, and may cause feelings of euphoria.

NIGHTSTAR LX™

Serious tools for serious apps

See if NightStar LX is right for you.
Download a **Free Trial** at

www.ccur.com/nightstar/LX

800-666-4544 • 954-974-1700
Email us at nightstar@ccur.com

NightStar LX is a product of Concurrent.

© 2005 Concurrent Computer Corporation

concurrent

MonarchComputer.com

AMD Solution Provider PLATINUM

Visit our website or Call 1-800-611-0875

PayPal

FREE! - Get Half-Life 2, a Patriot SD FlashCard Kit, and a LIVESTRONG™ Wristband when you upgrade to Athlon 64/FX/X2!

Purchase any AMD Athlon™ 64/FX/X2 bundle from Monarch and get Half-Life 2, a 128 MB Patriot SD Flash Card Kit w/USB Reader and a LIVESTRONG™ wristband for helping Monarch support the battle against cancer. Monarch Computer and you can make a difference.

PATRIOT MEMORY

FREE INSTALLATION SETUP & TESTING BY CERTIFIED TECHS

FREE TECH SUPPORT! Monarch makes it quick and easy to upgrade with FREE setup and testing on Motherboard Combos and $18.00 build fee on Barebones.

AMD Motherboard Combos

DFI K8T800PRO-ALF w/ AMD Sempron™ processor 2600+ (754 - 64 bit) — Only $129

Asus K8N-E Deluxe Mainboard with AMD Sempron™ processor 3400+ (754 - 64 bit) — Only $249

Asus A8N-E nForce4 Mainboard with AMD Athlon™ 64 processor 4000+ — Only $475

Abit AN8 SLi nForce4 with AMD Athlon™ 64 X2 processor 3800+ (Dual Core) — Only $539

Monarch A0110289 (H8DCE) Mainboard w/ AMD Opteron™ processor 246 — Only $585

Tyan S2895UA2NRF K8WE MB w/ NEW CPU *AMD Opteron™ processor 280 (Dual Core)* — Only $1795

Mainboard - Processors - Heatsink and Fan with Memory Options - FREE INSTALLATION AND TESTING
LASTEST BIOS loaded for easy upgrades - AMD Sempron™, Athlon™ 64, Athlon™64 FX, and Opteron™ Combos Available

AMD Barebone Systems

Go to www.monarchcomputer.com, select Barebones from the Menu. Choose AMD Sempron™, AMD Athlon™ 64 or AMD Opteron™. Then configure your barebones online or call 1-800-611-0875.

Enermax CS-10182-BA E-ATX Tower w/460W PS
Tyan S2875ANRF (Tiger K8W 8151) MB
AMD Opteron™ processor 242 1.6 GHz
Starting @ $589

Antec Sonata II Piano Black Quiet Mid Tower w/450W PS
Abit AN8 Ultra nF4
AMD Athlon™ 64 processor 3700+ (939 - 90nm)
Starting @ $509

Lian-Li PC-70 Aluminum Server Case w/550W PS
Tyan S2882G3NR-D (Thunder K8S)
AMD Opteron™ processor 270 (Dual Core)
Starting @ $1579

Apex, FoxConn Mid-Tower TU-150 Case w/400W PS
DFI K8T800PRO-ALF Motherboard
AMD Sempron™ processor 2600+ (754 - 64 bit)
Starting @ $217

***AMD Athlon 64 and Athlon 64 FX are the first Windows®-compatible 64-bit PC processor

www.monarchcomputer.com

AMD Price Drop!

AMD Athlon™ 64 Retail Box CPUs
AMD Athlon™ 64 3000+ (939) $109.00
AMD Athlon™ 64 3200+ (939) $152.00
AMD Athlon™ 64 3500+ (939) $201.00
AMD Athlon™ 64 3700+ (939) $233.00
AMD Athlon™ 64 3800+ (939) $282.00
AMD Athlon™ 64 4000+ (939) $334.00
AMD Athlon™ 64 FX-55 (939) $811.00
AMD Athlon™ 64 FX-57 (939) $1011.00

AMD Athlon™ 64 X2 Retail CPUs
AMD Athlon™ 64 X2 3800+ $322.00
AMD Athlon™ 64 X2 4200+ $400.00
AMD Athlon™ 64 X2 4400+ $497.00
AMD Athlon™ 64 X2 4600+ $630.00
AMD Athlon™ 64 X2 4800+ $787.00

AMD Opteron™ OEM CPUs
AMD Opteron™ 146 2.0GHz $121.00
AMD Opteron™ 148 2.2GHz $207.00
AMD Opteron™ 150 2.4GHz $311.00
AMD Opteron™ 244 1.8GHz $158.00
AMD Opteron™ 246 2.0GHz $203.00
AMD Opteron™ 248 2.2GHz $307.00
AMD Opteron™ 250 2.4GHz $441.00
AMD Opteron™ 252 2.6GHz $669.00
AMD Opteron™ 254 2.8GHz $825.00
AMD Opteron™ 848 2.2GHz $677.00
AMD Opteron™ 850 2.4GHz $847.00
AMD Opteron™ 852 2.6GHz $1130.00
AMD Opteron™ 854 2.8GHz $1469.00

AMD Opteron™ Dual Core OEM CPUs
AMD Opteron™ 265 1.8GHz $669.00
AMD Opteron™ 270 2.0GHz $825.00
AMD Opteron™ 275 2.2GHz $1019.00
AMD Opteron™ 280 2.4GHz $1260.00
AMD Opteron™ 865 1.8GHz $1130.00
AMD Opteron™ 870 2.0GHz $1469.00
AMD Opteron™ 875 2.2GHz $2085.00
AMD Opteron™ 880 2.4GHz $2570.00

The AMD Athlon™ 64 X2 dual-core processor provides the same level of system features customers have grown to expect with the AMD Athlon™ 64 product family: HyperTransport™ technology - Enhanced Virus Protection for Microsoft® Windows® XP- SP2 - Cool'n'Quiet™ technology

Components and Upgrades
1000s of In-Stock Components

100988
Antec SLK1650B Black Mini Tower w/350W PS
$70.00

150439
Western Digital 320GB SATA 8MB Cache 7200RPM (WD3200JD)
$135.00

100298
PC Power & Cooling Silencer 470W ATX 12V Power Supply (OEM)
$95.00

140226
2 GB DDR (2 pcs 1GB) PC-3200 (400) OCZ Dual Channel Platinum (OCZ4002048ELDCPE-K)
$259.00

190522
ATI (Connect3D) Radeon 9250 256MB DDR/8x-AGP TV-Out/DVI (Retail Box)
$56.00

150239
Western Digital 74 GB SATA 10K Raptor (WD740GD)
$175.00

150142
3Ware 9550SX-8LP Kit, 8 port SATA low profile PCI 64/133 RAID Controller
$515.00

3ware AUTHORIZED RESELLER

110292
Supermicro/Monarch A0110292 (H8DAR-E) GB-LAN/Video/SATA Reg-DDR/E-ATX Dual-Core 940
$339.00

Educational and Government POs Welcome.

GSA Schedule Contract GS-35F-0202P

We ship to the Continental U.S., Alaska, Hawaii, APOs, Puerto Rico, and Canada

Commercial leasing available for purchases as low as $1000.

Prices subject to change without notice. Monarch Computer not responsible for typographical errors.

AMD, the AMD Arrow logo, AMD Athlon, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. All brands and product names are trademarks or registered trademarks of their respective companies. AMD model numbers indicate relative software performance among this AMD processor family. *Enhanced Virus Protection (EVP) is only enabled by certain operating systems including the current versions of Microsoft® Windows®, Linux, Solaris, and BSD Unix. After properly installing the appropriate operating system release, users must enable the protection of their applications and associated files from buffer overrun attacks. Consult your OS documentation for information on enabling EVP. Contact your application software vendor for information regarding use of the application in conjunction with EVP. AMD and its partners strongly recommend that users continue to use a third party anti-virus software as part of their security strategy.

Monarch Has The LOWEST PRICES
Custom 64-Bit Servers, Workstations & Desktops
Available with AMD Dual-Core Technology!

AMD 64 Athlon X2 | AMD Sempron | AMD 64 Athlon FX | AMD 64 Athlon | AMD 64 Opteron

NEW! Monarch Server Overhaul!

Monarch's EMPRO™ line makes buying AMD Opteron™ Workstations and Servers Easier than ever before!

Take the Monarch Challenge: Give us 5 minutes and Save $100s if not $1000s! Come see our jaw-dropping prices! In most cases, we're cheaper than the competition BY AS MUCH AS 20%!

1. Visit our website and select "Servers".
2. Choose 1U, 2U, 3U or 4U and choose SATA or SCSI.
3. Custom configure your new server, or use one of our pre-configurations based on:

Price Point,	Application,	or General Use
Least Expensive	Fileserver	Max Storage RAID 5
Under $2000	Appserver	Standard RAID 5
Under $3000	Computational	Standard Mirror
Under $4000	AV Editing	Most Redundant
		Maximum Config.

4. Purchase your server or save your quote and lock in pricing for up to seven days. Only at MonarchComputer.com!

"BOTTOM LINE: MUST BUY"

LINUX MAGAZINE EMPEROR CLASS

"What's not to like? Monarch provides top parts, excellent customer service, and has earned the highest-level solutions provider status recognized by AMD and other key component vendors."

Jason Perlow
Linux Magazine
April 2005

For more information on the Empro™ Line Visit www.monarchcomputer.com/empro

"I bought the dream system with top components that were hard to find elsewhere. The guys at Monarch did a perfect job of building it, down to the smallest details. Wiring was perfect, position of hard drives within the chassis also perfect, etc. Even the packaging was superb. And every communication I've had with them was also perfect. Ordering from Monarch was a good move."

User: canbbb
ResellerRatings.com

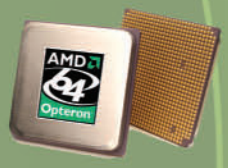

The AMD Opteron processor—built upon forward-thinking AMD64 technology—provides flexibility with a 1–8-way scalable design.

LINUX FRIENDLY!
CUSTOM AND PREBUILT PCS - Linux Preinstalled!
Fedora - RedHat - SuSe

<< LOOK FOR THIS SYMBOL ON LINUX COMPATIBLE MONARCH SYSTEMS!

SAVE MONEY SACRIFICE NOTHING!

9 MAXIMUM PC KICKASS | PC WORLD ★★★★☆ | CPU COMPUTER POWER USER | LINUX MAGAZINE EMPEROR CLASS | COMPUTER SHOPPER EDITORS' RATING 8.0

PRICEWATCH ★★★★☆ | RESELLERRATINGS.com | BizRate Overall Rating: | PriceGrabber.com ★★★★★ | Shopping.com Store rating:

empty string. We could solve the problem by modifying the definition of the People table, adding checks to ensure that the contents of each field is a non-empty string—but if we were to do this, Rails would show us the database error, complaining that we had violated an integrity constraint.

The solution is to modify the Person object so that it catches such errors, forcing the user to enter something in each field. We do this by modifying the Person class definition, located in app/models/person.rb. When we first open person.rb, we see that it is an unchanged subclass of ActiveRecord::Base:

```
class Person < ActiveRecord::Base
end
```

We can add one of the built-in Rails validators, statements that allow us to check the integrity of the data at the application level, before it ever gets to the database level. In this case, we use validates_presence_of, naming each of the fields from our table:

```
class Person < ActiveRecord::Base
    validates_presence_of :first_name, :last_name,
                          :email_address, :phone_number
end
```

With this in place—and without even having to restart the server—we can try adding another blank person. But now we find that Rails has stopped us, explaining the problem (for example, "Phone number can't be blank") at the top of the form and outlining each of the offending fields in red. With this validator in place, we can be sure that all of the rows in the People table will contain valid data.

When we go to /Appointments to add a new appointment, something seems suspicious even before we click on the create button at the bottom of the page: there isn't anywhere that we can enter the person with whom we are meeting! This will cause problems, as clicking on the create button quickly demonstrates; PostgreSQL returns an error, which Rails displays for all to see. Clearly, we need to solve this problem.

The problem is that the view for creating new instances of the Appointment class (that is, app/views/appointments/new.rhtml) is missing an HTML form element named appointment[person_id]. If new.rhtml were to include appointment[person_id], it would be submitted along with the rest of the elements of the form and inserted into the database.

The thing is, appointment[person_id] should be populated from the database. Assuming that we have a variable named @people available to us, we could add something like this to new.rhtml right before the call to submit_tag:

```
<b>Person:</b><br />
<select name="appointment[person_id]">
    <option value="">Select a person</option>
    <% @people.each do |person| %>
    <option value="<%= person.id %>">
    <%= person.first_name %>
    </option>
    <% end %>
</select><br />
```

The above RHTML code is similar to JSP and ASP in that it embeds Ruby code inside of an HTML document. Code surrounded by <% %> is executed in place, while code surrounded by <%= %> is replaced by its return value.

The above code thus defines an HTML form element named appointment[person_id]. It then creates an option with a blank value. Next, we get into a standard Ruby idiom, iterating over the elements of a list, using person as an iterator, pulling out person.id as the value and person.first_name as the text. In other words, we create a <select> list of the people in our People table.

But where does @people come from? We have to define it, but we can do that inside of the Appointments controller object, app/controllers/appointments_controller.rb. That file contains all of the methods the scaffolding system created for us. We merely have to add one line to the new method definition:

```
@people = Person.find_all
```

Now, we know that @people is a variable we're defining, and we know that Person is a subclass of ActiveRecord::Base that hooks us to the People table in our database. The find_all method returns all of the elements in the table.

Finally, we modify our data model class, appointment.rb, adding a validator to ensure that we will have nonblank values for each of the fields:

```
class Appointment < ActiveRecord::Base
    validates_presence_of :start_at, :end_at, :comment, :person_id
end
```

With all of this in place, we can begin to schedule appointments. Each appointment will be with a single person, and we can be sure that it will contain all of the data that we want. Moreover, we know that by the time PostgreSQL receives the data to be inserted, it will be valid.

Conclusion

Although constraints in our database ensure that the data will always be valid, we generally want to perform such validation at the application level. Unfortunately, doing so is tricky or time consuming in many languages. ActiveRecord, the object-relational mapper at the heart of Ruby on Rails, makes it relatively easy to ensure that your users never have to see a database error. It comes with a number of validators, as well as an infrastructure for creating custom ones. Moreover, it comes with a number of routines that let us describe the relationships among different tables. With some small modifications to the controllers, views and models, we are able to create a custom application with valid data quickly.

Resources for this article: www.linuxjournal.com/article/8580.

Reuven M. Lerner, a longtime Web/database consultant and developer, now is a graduate student in the Learning Sciences program at Northwestern University. His Weblog is at **altneuland.lerner.co.il**, and you can reach him at reuven@lerner.co.il.

Tools to build. Tools to tune. Tools to transform.

You're building software applications that will change the way the world works, plays and interacts. We're right there with you. With products that help your applications run better and faster. Compilers. Tuners. Threaders. Debuggers. And much more. Our products help you integrate software innovation with Intel platform technologies—from idea to optimization.

Intel® | software network

Check us out at www.intel.com/software/products

intel®

© 2005 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

amaroKing the Night Away

amaroK is a powerful music player and music management tool for Linux. BY MARCEL GAGNÉ

François, *oui*, it does look like a bit of a mess to me as well. I agree that it started out as a great idea, ripping all the music CDs in the restaurant to OGG format. The problem, *mon ami*, is that the mess has gone from the CD shelf to the hard drive of your Linux system, and it is getting only messier. *Non, mon ami*, I am not trying to make fun of you, but unless we start using some kind of media player software that make sense of this mess, your great idea will quickly become like most of your ideas.

Now, now, François, I am just kidding. You are a fantastic waiter and an unconventional thinker as well. You are welcome, *mon ami*. And I have just the software you need featured on tonight's menu. But quickly now, our guests will be here any moment and we need to prepare for them and select a wine. Too late! They are already here! Welcome, everyone, François, please head down to the wine cellar and bring back the 2002 Domaine Vincent Girardin Meursault Les Narvaux.

Today, *mes amis*, we are going to feature but one item on the menu—the breadth of its capabilities demands it. François has been trying to convert all his music to digital format but needs a combination media player and organizer—a digital jukebox, and I have just the thing for him. This amazing Linux software package is, in this humble chef's opinion, the best media player ever created, regardless of your operating system. It's called amaroK and it truly rocks.

amaroK's features are too numerous to list, but let me give you a sample of what the program offers. There's a powerful cover manager (downloads covers from Amazon.com), a context browser that keeps track of your favorite and most-listened-to songs, a skinnable interface, iPod support (other players work as well), great visualization tools (using libvisual) and more. There's even a lyric download feature so you can sing along with your favorite tunes without worrying about whether you are getting the words right. There is much more, and I will show you some of its capabilities in a moment.

The first step to getting amaroK running on your system (check first as some distributions come with it) is getting a copy (see the on-line Resources). There's really no need to build amaroK from scratch, as precompiled packages are available for an amazing number of distributions. Should you need to compile amaroK or choose to do so for the exercise, this is another example of a simple, extract-and-build five-step:

```
tar -xjvf amarok-1.3.1.tar.bz2
cd amarok-1.3.1
./configure --prefix=/usr
make
su -c "make install"
```

There are several options for building amaroK with a handful of audio engines from which to choose (by default, aRts and Helix are used), and as such, you may need other supporting packages (for example, gstreamer). The most likely one you may need, however, is taglib, a library used for reading and writing metadata and ID tags on MP3 and OGG files. Check your distribution CDs or visit the TagLib site (see Resources).

When you run amaroK the first time, you are presented with the aptly named First-Run Wizard (Figure 1). This is a simple three-step wizard that asks you to select an interface style, a folder where your songs are stored (this can be a high-level directory where you have access to the subdirectories) and the obligatory congratulations screen.

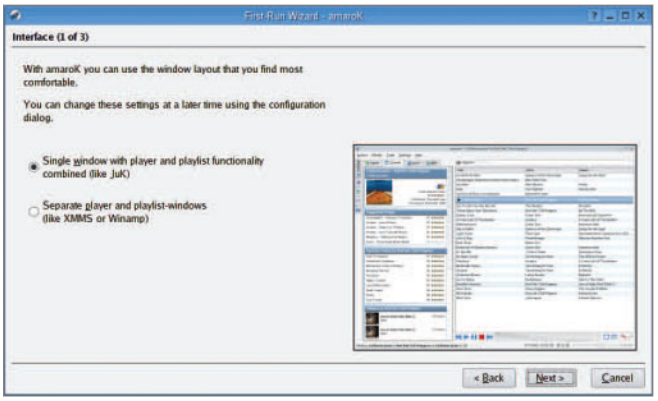

Figure 1. Setting up amaroK is easy with the First-Run Wizard guiding you along.

If you selected a folder to scan for songs, amaroK starts by scanning that folder and building a song collection. As part of this process, each song's ID tag is examined to build a list sorted by artist, song, album and so on. How long this takes depends on how many songs you have stored on your disk. To give you some idea, a progress status bar near the bottom of the amaroK window displays the percentage of completion (Figure 2).

Figure 2. As amaroK scans your folders, it automatically builds a collection of all your music.

Once that collection is done, you are ready to go, and you can start playing the song of your choice right away. Let's take a moment, however, to see how amaroK is laid out. The larger, right-hand pane contains your playlist. At the bottom of the

Boost Performance on Linux Clusters

Create, Analyze and Optimize High-Performance Applications on Clusters of Intel® Processor-based Systems

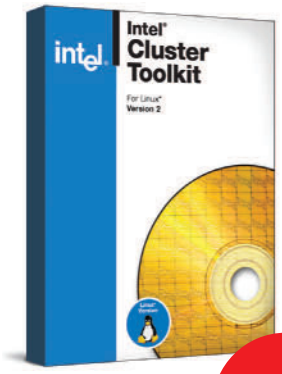

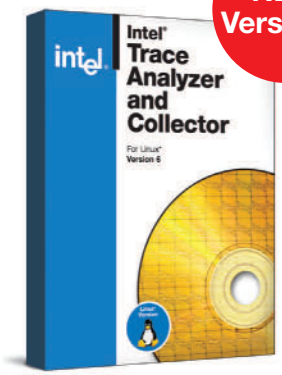

NEW Versions!

Introducing the Intel® Cluster Toolkit 2.0 for Linux

These tools address every stage of the development process, enabling you to achieve optimized performance for Intel® processor-based cluster systems.

Intel® Trace Analyzer and Collector 6.0

Optimized MPI Performance Analysis

Intel® Trace Analyzer and Collector bundle together graphical analysis and runtime event traces, providing maximum performance to cluster application developers.

- Navigate through the trace data at different levels of abstraction (cluster, node, and process) to address data scalability.
- View important aspects of application runtime behavior in detail and in the aggregate.
- Low-overhead, event-based tracing to profile the behavior of multithreaded applications with and without a Message-Passing Interface (MPI).

9.0 available NOW!

Intel® Fortran Compiler 9.0 for Linux

Compatible with Linux* development tools most programmers already use. A powerful combination of Compaq Visual Fortran (CVF) front-end and Intel® Fortran back-end, along with substantial CVF source code compatibility.

Intel® MPI Library for Linux

Use this high-performance message-passing library to develop applications that can run on multiple cluster fabric interconnects chosen by the user at run-time.

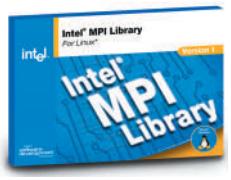

Programmer's Paradise®

To order or to request additional information

Call: 800-423-9990

Email: intel@programmers.com

FREE trials available at: programmersparadise.com/intel

YOU SAVE UP TO $52!	Paradise #	Retail	Discount
Intel® Cluster Toolkit for Linux 2.0	I23 0E70		CALL
Intel® Trace Analyzer and Collector 6.0	I23 0E71		CALL
Intel® Fortran Compiler 9.0 for Linux	I23 0BL3	$699.00	$646.99
Intel® MPI Library 1.0	I23 0BFP	$499.00	$474.99

© 2005 Intel Corporation, Intel, the Intel logo, Pentium, Itanium, Intel Xeon and VTune are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. *Other names and brands maybe claimed as the property of others.

playlist window are the controls to pause, play and jump to the next or previous song. There's also a volume control slider, a position slider (so you can move within the song itself) and a nice graphic sound analyzer display. Incidentally, if you click on the analyzer, it switches through a number of different display styles. Over on the left-hand side is a list of artists with some plus signs beside the artists' names. In both cases, there's a search box above the pane to let you find a particular tune quickly by typing in part of the name.

Let's get back to the playlist for a moment. Click on the plus sign beside an artist's name, and the entry expands to reveal the various albums by that artist, each of which further expands to list individual songs. Double-click on a song, and it appears in the playlist window and starts playing (Figure 3). I think, *mes amis*, I'll just double-click on Bachman Turner Overdrive's "You Ain't Seen Nothin' Yet" and crank up the volume a bit while François refills everyone's glasses.

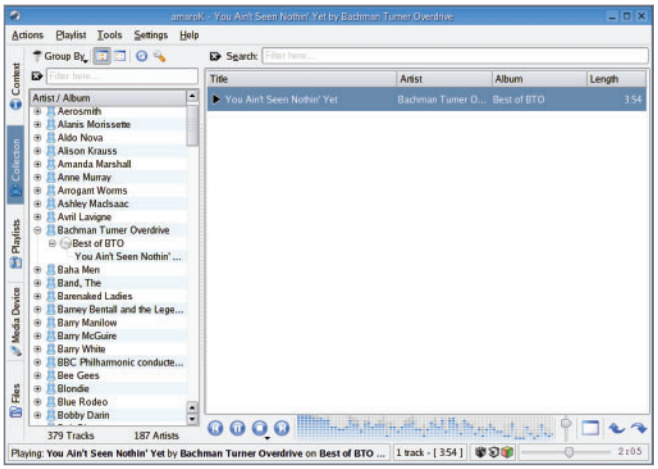

Figure 3. amaroK doing what it does best. In a hurry? Find a song, double-click, and sit back and listen.

If you want to queue up a number of songs, simply drag them from the list of songs into the playlist. As you build your playlist, you even can move the songs up and down in the list, changing the order of play as you go. So far, it sounds pretty good, but it gets a whole lot better. Before we get too far though, it's time to look at those tabs running down the left-hand side. When you are searching for songs, queuing them up for playing by dragging them into the playlist, the active tab is the Collection tab. For those impatient among you who just went ahead and started playing a song, you will have noticed that the left-hand tab and pane changed from Collection to Context. The context manager is one of the most powerful and useful features in amaroK and it deserves some explanation. When the context manager is active, there are four tabs along the top of the left-hand pane. They are labeled Home, Current, Lyrics and Artist.

The Home tab lists information about the most recently played tracks, your favorite tracks (based on how often you play them) and your newest tracks. The tab labeled Current displays information about the current track. This includes complete information (or as much as you have) on the current track and artist, the album cover (more on this shortly), the rating and when it was last and first played. It also lists your favorite songs by that artist and other albums you may have in your collection. You may even see one or more *suggested* songs listed. The Lyrics tab will query on-line lyrics servers to find the words for the current track (I personally love this feature since I like to accompany my music—enough smirking, François). Finally, the Artist tab will query the Wikipedia on-line encyclopedia to return the information relating to the current artist (Figure 4).

Figure 4. Find out all about your favorite artist from Wikipedia while you listen.

Since I mentioned the album cover in the context manager discussion, it's only fair that I go back and talk about this feature. Aside from all this great information about the track that's playing at any given moment, most people will tell you that it's kind of cool to have the album cover displayed as well. I personally don't want to go through the hassle of scanning my CD covers and storing all those pictures on my system, but amaroK makes this easy by downloading the cover art from Amazon.com. When a song is playing, the context manager displays information with either a default question mark cover or the actual cover (Figure 5). To download the cover for that particular song, right-click on the image and select Fetch from

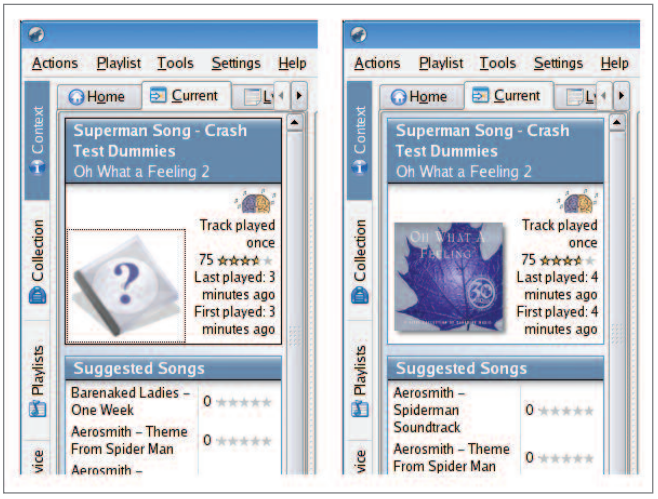

Figure 5. When songs are first loaded, there is no cover information. Right-clicking on the default question mark cover lets you download a cover from Amazon.com.

linux phone systems 101

lesson one

... choose the right team

The first step in any project is assembling the right team for the job. Let the experts at Telephonyware guide your Linux™ phone system project by helping you select the best hardware and software, and by providing the very best in service and support.

Take the guess work out of VoIP, choose a partner you can trust — Telephonyware.

lesson two

... get the right gear

For your VoIP project to be successful, you need the right gear! Let Telephonyware take the worry out of selecting the right hardware and software for the job.

We sell and support a full range of IP phones, analog and digital telephony cards, analog telephone adapters (ATAs), power over ethernet midspans and switches, and many more quality products. Our range is hand picked from the best manufacturers, and our helpful staff have used every product we sell.

lesson three

... put it all together

When it's time to turn plans into reality, Telephonyware is the right partner to take you from idea to completion. Our network of service partners, and excellent in–house support, give you the confidence you need, at a price you can afford.

Whether you're an experienced consultant deploying VoIP for your customers, a business replacing a phone system, or just looking for an IP phone or an ATA, Telephonyware will help you put it all together.

POLYCOM®

Telephonyware sells, supports and recommends the full range of Polycom IP phones.

Polycom's range of IP phones have been engineered to deliver a superb communications experience. They provide outstanding sound quality, advanced functionality, ease of use, simplified configuration and upgrades, and protection of your investment over time.

The Polycom VoIP portfolio includes the SoundPoint® IP family of desktop phones, an attendant console based on the SoundPoint IP 601 and Expansion Module, and the SoundStation® IP 4000 conference phone.

Both the SoundPoint range of desk phones, and the SoundStation range of conference phones seamlessly integrate with your IP PBX or softswitch application.

The phone's intuitive user interface offers dedicated, single button access to common telephony features. The high-resolution display delivers content for call information, multiple languages, directory access, system status and future applications.

The entire range of Polycom phones and accessories can be purchased by calling us, or directly though our web site.

Telephonyware is proud to be a Polycom Certified Channel Partner.

CERTIFIED RESELLER POLYCOM VOICE VoIP

... visit www.telephonyware.com/polycom for more info

TELEPHONY WARE

For online orders or more info, please visit us at **www.telephonyware.com/lj**

Call us on **(866) 864-2304** or write to **sales@telephonyware.com**

Polycom and SoundPoint are registered trademarks of Polycom, Inc. All other trademarks are the property of their respective owners.

No true entertainment system is complete without a great light show to go with it.

Amazon.com from the menu that appears.

This is a great way to import cover art for the occasional track, but you may have already had hundreds of songs stored on your PC when you decided to use amaroK, and getting all those covers individually can take an amazing amount of time. Luckily, there is a better way. Just use amaroK's cover manager. Start by clicking Tools, then select Cover Manager. The Cover Manager window appears with a list of all the covers for which you have albums identified (a song without an album title in its meta tags won't show up here). You'll probably see a whole slew of albums with the default question mark cover. Now, look up in the top right-hand corner of the Cover Manager and you'll see a button labeled Fetch Missing Covers. Click that button, sit back and wait while amaroK does the rest (Figure 6).

Figure 6. amaroK's Cover Manager makes downloading cover art a snap.

By now, you might be starting to believe that amaroK is as amazing as François and I do, *non*?

Need more convincing? No true entertainment system is complete without a great light show to go with it. Click Tools and select Visualizations. Not getting the right sound for your room or music style? You might need to change the levels using the built-in equalizer. Click Tools and select Equalizer. When the Equalizer window appears (Figure 7), the equalizer itself initially will be deactivated. Click the Enable Equalizer check box, and adjust it to your liking. The change in playback will accompany your changes. What's particularly interesting here is that there is a drop-down box near the top of the window labeled Presets. Click here and you'll find a number of preset levels suited to different musical styles such as Club, Large Hall, Pop, Rock, Reggae, Techno and several others.

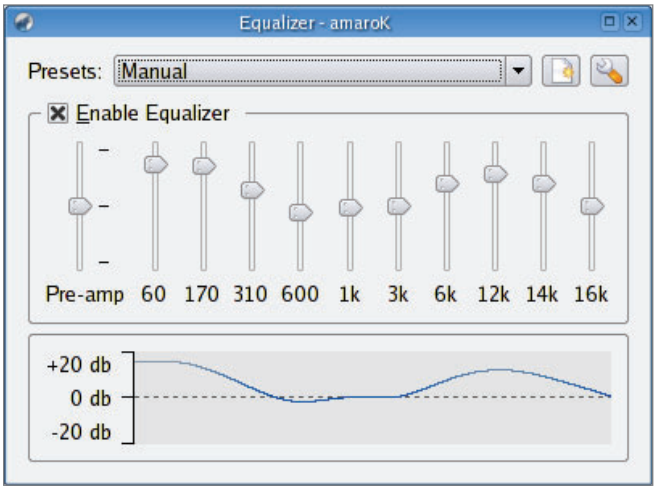

Figure 7. The equalizer can be set manually, but includes many presets for different musical styles.

amaroK also can play songs randomly, repeat your playlists for endless music, save playlists and then drag the whole thing down into your mounted iPod. Simply click the Media Player tab. Click the Playlists tab for even more multimedia fun. This is where you manipulate your various playlists, download and listen to podcasts or listen to Internet radio streams (amaroK is already configured with a number of these stations). With this media player on your Linux system, the entertainment never ends.

Can it be possible? The clock on the wall must be playing a joke on us, saying it is near closing time. With the music playing and François ready to refill your glasses, surely we can delay our parting a little longer. We'll drag a few more songs into the playlist, turn up the volume just a little higher, and maybe see if we can't find some truly decadent *Gâteau au chocolat* to finish off the evening. Please raise your glasses, *mes amis*, and let us all drink to one another's health. *A votre santé! Bon appétit!*

Resources for this article: www.linuxjournal.com/article/8582.

Marcel Gagné is an award-winning writer living in Mississauga, Ontario. He is the author of *Moving to Linux: Kiss the Blue Screen of Death Goodbye!* 2nd edition (ISBN 0-321-35640-3), his fourth book from Addison-Wesley. He also makes regular television appearances as Call for Help's Linux guy. Marcel also is a pilot and a past Top-40 disc jockey. He writes science fiction and fantasy and folds a mean Origami T-Rex. He can be reached via e-mail at mggagne@salmar.com. You can discover a lot of other things (including great Wine links) from his Web site at www.marcelgagne.com.

MORE SPACE. LESS MONEY.

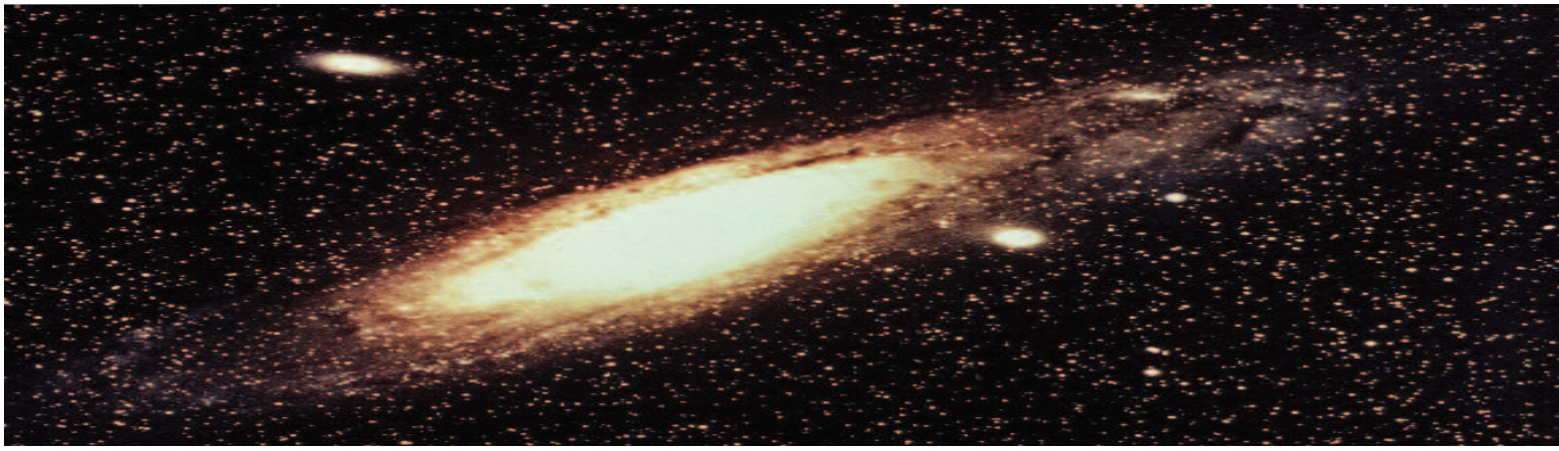

UNLIMITED AFFORDABLE NETWORK STORAGE

Everybody needs more space. And they need to spend less money. What if you can both have more space and spend less money?

What if you could put 7½ terabytes in only 3 rack units? What if that 7½ terabytes cost less than $10,000? Including the SATA disk drives. Imagine if you could glue it all together with a RAID appliance into one system. What if you could add as much storage as you wanted, one shelf at a time, and never have to `fork-lift' anything?

Coraid's new SATA EtherDrive Storage allows you to do just that. Using industry standard SATA disk drives, EtherDrive Storage connects disks directly to your Ethernet network. Each disk appears as a local drive to any Linux, FreeBSD or Solaris system using our open ATA-over-Ethernet (AoE) protocol. Since the disks just appear as local drives you already know how to use them.

The EtherDrive® SATA Storage Shelf is a 3U rack-mount network appliance that contains 15 SATA drive slots. Its triple redundant power supply protects you from your most likely failure. Its dual Gb Ethernet interfaces allow your data to go fast; 200MB per second. And at a very affordable price. List price for the EtherDrive Storage Shelf, without disks, is only $3,995.

Our companion product, the RAIDBlade RAID controller, allows a virtually unlimited number of Storage Shelves to be combined into a set of logical AoE storage devices.

Now you can have unlimited storage at a very affordable price. For complete information, visit our website at www.coraid.com, or call, toll-free, 1-877-548-7200. And we'll show how we've made network storage so affordable, you can have all the space you want.

www.coraid.com
info@coraid.com
1.706.548.7200

CORAID

Get Started with Redirection

Dave Taylor, author of *Wicked Cool Shell Scripts*, begins a new series on Linux shell scripting in this issue.

BY DAVE TAYLOR

If you're reading this publication, you already know that Linux is one of the most powerful and versatile operating systems available today. If you're an old-timer like me, you also know all about the command line and the geeky retro joy that typing commands rather than clicking icons offers the diligent user. Nowadays, though, the graphical interface layered atop Linux is so well designed that—though I find it a bit baffling—plenty of Linux users never go near the command line.

That's too bad. The command line is tremendously powerful, and the underlying metaphor of commands being strung together in pipes to create custom command sequences means that Linux actually offers millions of unique ways to work with the system. But, yes, there's a definite learning curve to overcome.

More than just the command line, though, it turns out that the shell offers a simple and surprisingly powerful programming environment through what we call shell script programming. In UNIX parlance, a shell is a command-line interface or CLI. Either way, it's the program that receives the commands you type in and actually does whatever it is you requested. String a bunch of these commands together, put them in a file and you have a shell script—simple and straightforward.

That's what I'm going to address in this new column here at *Linux Journal*, and fair warning for those über-geeks in the crowd, I'm going to go slow and make sure we cover all the basic concepts before we move into complex scripting tricks and techniques.

To start, let me briefly introduce myself. I first logged in to a BSD UNIX system way back in 1980 and have been involved with UNIX, and then Linux systems, ever since. I worked with the Open Software Foundation, helped manage the Usenet hierarchy, was one of the postmasters at hplabs back in the old UUCP days and am pretty well known as the author of The Elm Mail System. I've written 19 books, notably including *Teach Yourself Unix in 24 Hours* and the best-selling *Wicked Cool Shell Scripts*. I've contributed software to a variety of UNIX and Linux distros, including BSD 4.4 back when that was released, and I still have an open terminal window on my computer regardless of what I'm working on. I'm hooked on the command line, what can I say?

Redirecting Input and Output

To get started, let's talk about one of the most important concepts of the Linux command line: standard input and output. When you run a program like `ls` to list files or `date` to see the date and time (sadly, the latter command doesn't help you gain a social life. If only it were so easy!), it turns out that the program actually has an input channel and two output channels. For these commands, the input channel is ignored because they don't actually read input from what's called the input stream, but they do have both an output and error output stream that are utilized. These three streams are called standard input (or stdin), standard output (or stdout) and standard error (or stderr). Why is this important? Because you can redirect any of them to come from a file or to go to a file—for any Linux command.

Let's say that you want to create a new file called rightnow, and you want it to contain the current date and time. Here's how that'd look on the command line:

```
date > rightnow
```

Easy enough. An important warning, however, is that if the output file you specify already exists, by default Linux just silently overwrites it, not infrequently leading to curses, great frustration and unhappy users. Be careful (or read up in your favorite command shell's man page about noclobber).

Let's say you want to save the date twice in the file. Now, instead of creating a new file, it's time to add the new content to the existing contents of the file. This is done thusly:

```
date >> rightnow
```

Check the file now and you'll see two time/date stamps, a few seconds apart.

Let's add another useful command to our list, wc, which counts characters, words and lines in either a specified file or in stdin (the standard input stream). First, how many characters, words and lines are in the standard output of the date command?

```
$ date > test
$ wc test
       1       6      29 test
```

Typical cryptic Linux output: the first value is the number of lines, the second the number of words and the third the number of characters. Let's try a variation on this too:

```
$ wc < test
       1       6      29
```

Notice this time that rather than having the wc command open up a file we've specified by name, we're using a

redirection to replace stdin with the contents of the specified file. That's why the wc output doesn't show the filename; it doesn't know that the input is from a file.

Let's consider one more file redirection before we wrap up this quick tour. We've seen > and >> and <. What do you think happens if you use << as a file redirection? Ah, well, that's a tricky one because it doesn't append anything, it lets you simulate a file redirection without actually having a file involved. In fact, << is known as a here document, because when used in the standard form of << EOF, it is read as "read until you reach 'here'" (the EOF sequence). This'll make more sense with an example:

```
$ wc << EOF
> this is a simple test and should
> show you how many lines, words
> and characters are in this little
> input sequence.
> EOF
       4       21      114
```

Now you can see where the output of wc is starting to make sense: four lines, 21 words and 114 characters. Count it for yourself! Also, notice that the > symbol at the beginning of the lines is automatically added by the shell as a continuation character to let you know that more input is expected. Once at the end of the here document, the sequence EOF appears, the input stream is fed to the specified command and wc dutifully counts lines, words and characters.

That should get us started with the basics this month. Next month, we'll explore how you can create pipelines of commands where the output of one command is the input of the next, then begin to talk about my long-term shell script programming project for this column: a rudimentary blackjack game.

Dave Taylor is a 25-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.

c-tree

Need a sharper development tool for your application's database?

SQL is only one of our options...

AMD 64 Opteron

C-TREE PLUS® DATABASE TECHNOLOGY OPENS UP YOUR OPTIONS

SQL offers a convenient and easy-to-use database interface. ISAM provides powerful performance with precision indexing control in a small footprint. With c-tree Plus you can simultaneously enjoy BOTH! Superior ISAM indexing technology PLUS an industry-standard SQL interface provide blazing fast data management for every environment. Break the limitations of a single solution and open up your database options. Experience the benefits c-tree Plus can deliver to your application!

FEATURED HIGHLIGHTS

- Fast, reliable, and portable
- Low deployment cost
- No DBA required
- Professional technical support
- Source code
- 64-bit support
- 16-exabyte file support
- Memory files
- Embeddable database
- Full OLTP support

CUSTOMER TESTIMONIAL

"We have reviewed Oracle and some of the other big relational databases and chose FairCom for our database development needs. With c-tree Plus, we see transactional volume that is 8 to 10 times faster than what we can get with other databases. I have been using c-tree based solutions since the 80's and highly recommend it..."

Visit our Web site for more testimonials about c-tree!

FairCom®
Database your way.

See for yourself — download c-tree Plus® Today!

Go to www.faircom.com/go/open for a FREE evaluation of c-tree Plus!

Other company and product names are registered trademarks or trademarks of their respective owners. © 2005 FairCom Corporation

Single Sign-On and the Corporate Directory, Part I

Author Ti Leggett presents the first in a series of articles focused on building a secure corporate directory, including support for single sign-on that's scalable up to thousands of users. BY TI LEGGETT

So you want a corporate directory, but you don't have a corporate budget. You want to reap the benefits of single sign-on, the ease of administration for yourself and the ease of use for your users. If you want all this, plus a secure and unified authorization and identity management system, read on. I'll start you down the path to sysadmin nirvana. In this series of articles, I'll show you how to build on pieces you may already have in place, add new pieces and make them all work together. Everything from the authentication servers, to mail delivery, to client integration (including Windows and OS X) will be discussed. We have a lot to cover, so let's get started!

Using Previous Building Blocks

We use MIT Kerberos V v1.4.1 and OpenLDAP v2.1.30 running on Gentoo Linux as our authentication and identity management systems, respectively. I assume you have three servers: kdc.example.com, ldap.example.com and mail.example.com. Before we go any further, you should first read the *Linux Journal* articles "Centralized Authentication with Kerberos 5, Part I" and "OpenLDAP Everywhere" (see the on-line Resources). We build on where those articles leave off, but keep in mind that our Kerberos realm will be CI.EXAMPLE.COM, and our base DN will be o=ci,dc=example,dc=com. Also, all of the configuration files referred to in this article are available from the on-line Resources.

Setting Up an SSL Certificate Authority (CA)

This section is optional reading but is highly recommended for sites that will have many servers using SSL. Each server can self-sign its own certificate, but you lose unity and some of the power of running your own CA. If you're interested in the details of OpenSSL, I highly recommend the book *Network Security with OpenSSL*.

We start by choosing /etc/ssl/example.com as the base directory to store all the signed certificates, certificate revocation lists (CRLs) and accounting information. Once that directory is created, we then create the directories certs, crl, newcerts and private underneath the base. We create an empty file /etc/ssl/example.com/index.txt, and then create a file /etc/ssl/example.com/serial:

```
# touch /etc/ssl/example.com/index.txt
# echo '01' > /etc/ssl/example.com/serial
```

Finally, we create the CA's OpenSSL configuration file, /etc/ssl/example.com/ca-ssl.cnf.

To create a self-signed CA certificate, we must do the following as the user who owns the /etc/ssl/example.com directory and its children, which is probably root:

```
# export OPENSSL_CONF=/etc/ssl/example.com/ca-ssl.cnf
# openssl req -x509 -days 3650 -newkey rsa \
 -out /etc/ssl/example.com/ci-cert.pem -outform PEM
# cp /etc/ssl/example.com/ci-cert.pem /etc/ssl/certs
# /usr/bin/c_rehash /etc/ssl/certs
```

For more details on the openssl req command, view the req(1) man page.

It is important to keep the passphrase for the CA key in a very safe place, because if the CA private key is compromised, all previously signed certs cannot be trusted. It is also important to keep the actual CA machine and access to it secure. How secure you keep the machine is up to you and your actual security needs, but if unauthorized users gain physical or network access, they have access to the CA private key. As I mentioned above, compromise of the CA private key compromises the entire chain of trust, making all signed certificates suspect and untrustworthy. Some suggest that the CA machine be physically secured with no network access. In order to sign certificates in this environment, you use registration authorities (RAs) to receive certificate signing requests (CSRs). The CSRs are then transferred to some secure portable media that is taken to the CA where the CSRs are signed, and the certificates written back to the portable media to be placed back on the RA for the end user to retrieve. If you think your needs might require this, the OpenCA Project was designed with this type of security in mind. It also has support for storage of signed certificates in LDAP.

We have created an OpenSSL configuration file for our CA, but that describes only how to request and sign exactly one certificate. We still need to create an OpenSSL configuration to use from now on to request normal host and user certificates: /etc/ssl/example.conf/ssl.cnf. The client configuration is a little more complex than the CA's because more variations can occur for client certificates.

Now that we have a client configuration file, let's generate a host certificate for the LDAP server. Generating a CSR can be done as a normal user:

```
# export OPENSSL_CONF=/etc/ssl/example.com/ssl.cnf
# openssl req -new -nodes -keyout ldap-key.pem \
```

```
-out ldap-req.pem
```

The openssl options used are much the same as those used for generating the CA CSR. The only new option is the -nodes option, which creates an unencrypted private key.

Our next step is to have the CSR signed by the CA in order to get the public certificate. This, again, needs to be done as root:

```
# export OPENSSL_CONF=/etc/ssl/example.com/ssl.cnf
# openssl ca -policy policy_anything -out \
 ldap-cert.pem -in ldap-req.pem
```

At this point, we have three files: ldap-cert.pem, the public certificate; ldap-key.pem, the private key; and ldap-req.pem, the CSR. The CSR can be thrown away once the certificate has been signed by the CA. Again, protecting the private key is important, especially because it is not encrypted. It probably should be owned by root and have permissions 0400.

Securing LDAP

Even though passwords aren't stored in the LDAP directory, a lot of sensitive information is. Your users probably don't want the whole Internet to know their phone numbers, e-mail addresses or employee IDs. Once you've read "OpenLDAP Everywhere" and have a working LDAP server, you need to secure the information transportation and access to the directory.

The first step is to secure the data transport using OpenSSL. First, let's copy our certificate and key we signed previously to /etc/openldap/ssl/slapd-cert.pem and /etc/openldap/ssl/slapd-key.pem, respectively. We need to provide five options in slapd.conf: TLSCipherSuite (optional), TLSCACertificatePath, TLSCertificateFile, TLSCertificateKeyFile and TLSVerifyClient. The slapd.conf(5) man page has good definitions of these options.

Having secured the data on the wire, we now secure authentication using the Kerberos KDC. OpenLDAP is Kerberized and uses SASL for authentication negotiation. We first must tell slapd how to find its Kerberos keytab file. We do this by editing /etc/conf.d/slapd or by defining KRB5_KTNAME prior to starting slapd in its init script. Two options in slapd.conf also must be defined: sasl-secprops and sasl-regexp.

Right now, TLS and SASL can be used but aren't required. Two more options in slapd.conf, security and allow, are used to specify the security methods and encryption strength needed for certain operations to take place. And, be sure to set up access control lists (ACLs) properly—refer to slapd.access(5).

Securely Replicating Kerberos

We start by replicating our Kerberos database from kdc.example.com to ldap.example.com, so that if kdc.example.com fails, ldap.example.com will pick up the slack. One important fact to remember is that only one kadmin server can be on the network for a realm at any time. Otherwise, there is no authoritative source for updates to the database. Kerberos comes with kprop and kpropd to propagate the Kerberos database securely. First we must identify kpropd as a known service. Add the following to /etc/services:

```
krb5_prop               754/tcp
```

We need to define an ACL file, /etc/krb5kdc/kpropd.acl, that tells kpropd what hosts are allowed to propagate. All that is really needed in this file is the master KDC's principal name, but it doesn't hurt to have all KDCs in here so that if a failure occurs, we can choose a new master, start the kadmin service on it and propagate from that host to the other slaves.

We now create an xinetd service definition, /etc/xinetd.d/kpropd, on our slaves; (re)start xinetd; dump the database on kdc.example.com; and propagate it to the slaves so they have an initial configuration:

```
# /usr/sbin/kdb5_util dump /etc/krb5kdc/slavedump
# /usr/sbin/kprop -f /etc/krb5kdc/slavedump \
 ldap.example.com
```

Finally, we create a stash file on each slave using the master key defined when setting up kdc.example.com's database, and then start the kdc service:

```
# /usr/sbin/kdb5_util stash
# /etc/init.d/mit-krb5kdc start
```

To propagate out the KDC database periodically, we define a cron job on kdc.example.com. Thanks to Jason Garman and the O'Reilly book *Kerberos: The Definitive Guide* for the original cron job.

A sensible time frame to run this script is hourly or from /etc/cron.hourly. Our Kerberos database is now being replicated securely from the master to any number of slaves. If the master fails, we have a way to switch to a slave machine quickly and with minimal data loss, if any. Now that we're propagating Kerberos changes, we can add the slave server to the krb5.conf file as a valid KDC.

Securely Replicating OpenLDAP

Enough critical information will be stored in your LDAP directory that you probably don't want a single point of failure. After all, if your LDAP directory is unavailable, your users won't be able to log in, check e-mail or do numerous other daily tasks. Replicating your LDAP directory helps ensure there is no single point of failure.

Let's replicate the LDAP directory from ldap.example.com to kdc.example.com. OpenLDAP has a dæmon called slurpd that is responsible for this. Unfortunately, slurpd has no configuration directive telling it which Kerberos keytab to use, so there's a bit of work required. First, we edit slapd.conf on ldap.example.com, adding the options replogfile and replica, and then we restart slapd.

We need to create a Kerberos ldap service principal and SSL certificate and key for kdc.example.com, as we did for ldap.example.com. We also must create a slapd.conf file for kdc.example.com. This file is almost identical to the one on ldap.example.com, with a few key differences. For the same reason we have only one Kerberos admin server, we want only one LDAP directory being updated and changed. The only user who should be able to write to the slaves' directory should be uid=host/ldap.example.com,cn=GSSAPI,cn=auth or the Kerberos principal of the master, so our ACLs on the slaves are

Replicating your LDAP directory helps ensure there is no single point of failure.

much more restrictive. Also, slapd needs to know who will be sending updates via slurp as defined by the updatedn and updateref options.

Now we switch our focus back to ldap.example.com for a bit. We need to create an /etc/conf.d/slurpd or make sure that KRB5CCNAME is set before slurpd is started from the init script.

Next, we get some initial Kerberos credentials:

```
# KRB5CCNAME=/var/run/slurpd.krb5cache /usr/bin/kinit -k
```

And then we dump the directory to a file:

```
ldap# /etc/init.d/slapd stop
ldap# /usr/sbin/slapcat -l /tmp/slavedump.ldif
ldap# /etc/init.d/slurpd start
```

Because slurpd transfers only changes in the master directory, we need to populate the slave directory with the current state of the master directory. We do this by copying a dump of the master we created above, /tmp/slavedump.ldif, to kdc.example.com and import the dumped directory and start slapd:

```
kdc# /usr/sbin/slapadd -l slavedump.ldif
kdc# /etc/init.d/slapd start
ldap# /etc/init.d/slapd start
```

We need to test that the slave has a sane directory:

```
# ldapsearch -H ldap://kdc.example.com -ZZ
```

To test that replication is happening, we can make a modification or addition to the directory on ldap.example.com and then search on kdc.example.com to make sure that change propagated.

Once we've verified that slurpd is working, we create a cron job on ldap.example.com to keep the credentials from expiring. The default time limit for credential validity is ten hours, so if we define a cron job to run every eight hours, we should be safe.

Last, we add kdc.example.com into our rotation of valid LDAP servers for nss_ldap. That is, we append kdc.example.com to the list of servers specified by the host option in /etc/ldap.conf.

Configuring the Postfix MTA

We'll be using the Postfix mail transport agent (MTA) v2.1.5. Postfix has well-established support for SASL authentication as well as LDAP support for features such as aliases. Because configuring Postfix from the ground up is beyond the scope of this article, we deal with how to enable Postfix to use SASL and TLS. For information on setting up Postfix, see Resources.

Postfix has two main configuration files, /etc/postfix/main.cf and /etc/postfix/master.cf. The main.cf file is primarily responsible for how to accept incoming mail, and master.cf is primarily responsible for defining mail delivery agents.

An example main.cf is included in the on-line Resources, but to understand the directives in this file fully, you should refer to the Postfix documentation and Web site.

Three main directives define how our SMTP server interacts with other SMTP servers: smtp_sasl_auth_enable, smtp_use_tls and smtp_tls_note_starttls. If your SMTP server will be exposed to the Internet at large, you should set these as flexibly as possible to ensure all other SMTP servers can talk to yours. If it's an internal-only SMTP server, however, you can make it more secure by strengthening these directives.

The more interesting part is how we specify the way our users and machines connect to our MTA to send mail. A few more directives are of concern here: smtpd_sasl_auth_enable, smtpd_sasl_security_options, smtpd_sasl_tls_security_options, smtpd_use_tls, smtpd_tls_cert_file, smtpd_tls_key_file and smtpd_tls_auth_only.

If you'll be using IMAP for mail delivery, make sure to set the mailbox_transport directive and the smtp and cyrus transports mechanism in master.cf.

Like OpenLDAP, Postfix is Kerberized, uses SASL for authentication negotiation and can use SSL to secure the data transport. To secure Postfix and configure it to use SASL, we need to do a few tasks in addition to modifying main.cf. First we create an SSL certificate/key pair and place the two parts in /etc/ssl/postfix/smtp-cert.pem and /etc/ssl/postfix/smtp-key.pem, making sure that they're owned by the user postfix and group mail, and that the key is readable only by user postfix. Next, we create a host principal for mail.example.com and save it to the normal place. We also create a service principal, smtp/mail.example.com@CI.EXAMPLE.COM and save it to /etc/postfix/smtp.keytab. This file should be owned by root and have the same permissions as the smtp-key.pem file. In addition, we create a SASL configuration file named /etc/sasl2/smtpd.conf and also edit /etc/conf.d/saslauthd. Postfix uses the saslauthd dæmon to get information about authentication mechanisms, and these two files tell SASL how to check passwords, what mechanisms are supported and the minimum security layer to use. Values for minimum_layer are equivalent to the security strength factors (SSFs) in OpenLDAP. Finally, we tell Postfix where its Kerberos keytab file is by creating /etc/conf.d/postfix or by making sure the KRB5_KTNAME environment variable is set in the init script prior to starting Postfix. Once all these tasks have been done, we can start the saslauthd and Postfix init scripts.

LDAP is useful not only for identity management and authorization but also for storing alias maps for Postfix. It's simple to use and maintain, and it removes the need to rebuild the alias database every time there is a change to it. The first step is to make our directory aware that we want to store alias maps in it. We do this by adding the misc.schema to the slapd configuration. Next, we create a branch in the directory for

AMD64 technology enables you to run all of your existing applications today while providing a seamless transition to 64-bit computing as your needs evolve.

Server System Comparison	AMD Opteron™	Intel Xeon MP	Intel Itanium 2
Modular, glueless scalability	Yes	Requires Northbridge	Requires Northbridge
SMP Capabilities	Up to 8-way	Up to 4-way	Up to 4-way
Direct Connect Architecture	Yes	No	No
Dual-Core technology	Yes	No	No
High Performance 32-bit and 64-bit computing	AMD64	EM64T	No
HyperTransport™ technology	Yes	No	No
Integrated DDR memory controller	Yes	No	No
Front Side Bus frequency	1.4 – 2.6GHz	667MHz	400MHz
Front Side Bus bandwidth	11.2 – 20.8GB/s	10.6GB/s	6.4GB/s
Maximum Inter-processor bandwidth	8.0GB/s	10.6GB/s	6.4GB/s
Memory support	DDR200/266/333/400	DDR266/333/DDR2-400	DDR200
Memory Bandwidth 2P System	12.8GB/s	12.8GB/s	6.4GB/s
Memory Bandwidth 4P System	25.6GB/s	12.8GB/s	6.4GB/s
L1 cache size (max.)	64KB (Data) + 64KB (Instruction) per core	16KB + 12k mop	32KB
L2 cache size (max.)	1MB per core	1MB	256KB
L3 cache size (max.)	N/A	8MB	9MB
Maximum I/O bandwidth 2P System	24.0GB/s	14.0GB/s	6.4GB/s
Maximum I/O bandwidth 4P System	32.0GB/s	14.0GB/s	6.4GB/s
SIMD Instruction Set Support	SSE, SSE2, SSE3	SSE, SSE2, SSE3	N/A

All New Powerful 64bits **Dual Core** Server Line

Friendly Server Specialists Complete Solution Provider

Flexibility and Variety : Customize platforms with leading-edge products

ZT Dual-Core Opteron 1U Server A8109

- **2 x Dual-Core AMD Opteron™ Processor 280 (2.40GHz)**
- AMD® 8132 + 8111 Chipset Server Board
- 2GB ECC Registered DDR 400MHz SDRAM (Supports up to 16 GB)
- 4 x Seagate® 300GB SATA Hard Drive (Total 1.2TB)
- 4 x 1" Hot-swap SATA Drive Bays
- 4 Channel SATA RAID Controller (RAID 0, 1, 5, 10 ,JBOD Support)
- Slim 24x CD ROM & Slim Floppy Drive
- 2 x 10/100/1000 Gigabit LAN
- 1U Rackmount Chassis w/500W Cold-Swap Power Supply
- Super Doctor III Server Management Software
- 3-Year Limited Warranty

$5,999

ZT Dual-Core Opteron 4U Server A8110

- **2 x Dual-Core AMD Opteron™ Processor 280 (2.40GHz)**
- AMD® 8131 + 8111 Chipset Server Board
- 4GB ECC Registered DDR 400MHz SDRAM (Supports up to 16 GB)
- 4 x Seagate® 400GB SATA Hard Drive (Total 1.6TB Storage)
- 8 x 1" Hot-swap SATA Drive Bays
- 8 Channel SATA RAID Controller (RAID 0, 1,5,10 ,JBOD Support)
- CD-RW / DVD Combo and Floppy drive
- 2 x 10/100/1000 Gigabit LAN
- 4U Rackmount Chassis w/650W Redundant-Cooling Power Supply
- Super Doctor III Server Management Software
- 3-Year Limited Warranty

$6,699

ZT Quad Opteron Server A8111

- **4 x Dual-Core AMD Opteron™ 875 Processor (Total 17.6GHz of CPU Power)**
- AMD® 8131 + 8111 Chipset Server Board
- 8GB ECC Registered DDR 400MHz SDRAM (Supports up to 64GB)
- 2 x 36GB 10,000rpm SATA Hard Drive (Raid 1 for O/S)
- 8 x 400GB 7,200rpm SATA Hard Drive (Raid 5 for Data)
- 16 x 1" Hot-swap SATA Drive Bays
- 16 Channel SATA RAID Controller (RAID 0, 1,5,10 ,JBOD Support)
- DVD±RW Burner and Floppy Drive
- 2 x 10/100/1000 Gigabit LAN
- 4U Rackmount Chassis w/1350W Hot-swap Redundant P/S
- 3-Year Limited Warranty

$16,999

Free Shipping for Holiday !

Find out how ZT System Specialist can help maximize your Business Solution

Go to **ztgroup.com/go/linuxjournal**

Call **866- ZTGROUP**(866-984-7687)

Promote code: lj1205

ZT GROUP

Purchaser is responsible for all freight costs on all return of returns of merchandise. Full credit will not be given for incomplete or damaged returns. Absolutely no refunds for merchandise returned after 30 days. All prices and configurations are subject to change without notice and obligation. Opened software is non-refundable. All returns have to be accompanied with an RMA number and must be in re-sellable condition including all original packaging. System's picture may include some equipment and/or accessories, Which are not standard features. Not responsible for errors in typography and/or photography. All rights reserved. All brands and product names, trademarks or registered trademarks are property of their respective companies.AMD, the AMD Arrow logo, AMD Athlon, Cool'n'Quiet and combinations thereof, are trademarks of Advanced Micro Devices, Inc. Windows is a registered trademark of Microsoft Corporation in the United States and/or other jurisdictions. Hyper Transport is a licensed trademark of the HyperTransport Technology Consortium. Other Product and company names used in this publication are for identification purposes only and may be trademarks of their respective companies.

PGI® Compilers are building the 64-bit applications infrastructure.

PGI Fortran, C and C++ compilers deliver world-class performance on a wide spectrum of 64-bit scientific and engineering applications. With PGI you get an easy-to-use integrated suite of dual-core and MPI-capable compilers, debugger, and profiler to simplify porting and tuning of 64-bit applications for AMD64 and EM64T processor-based workstations, servers and clusters. With comprehensive cross-platform support for Linux and 64-bit Windows operating systems on both Intel and AMD processors, PGI delivers a uniform development environment across your key target systems. The leading independent software vendors in structural analysis, computational chemistry, computational fluid dynamics, and automotive crash testing have chosen PGI compilers and tools to build and optimize their 64-bit applications.

Visit ***www.pgroup.com*** to learn what PGI Compilers and Tools can do for you.

The Portland Group™

www.pgroup.com ++ 01 (503) 682-2806

The registered trademarks and marks are the property of their respective owners.

STMicroelectronics

the aliases. We'll use ou=aliases,o=ci,dc=example,dc=com. The last piece is to tell Postfix to use LDAP as a source for aliases by adding ldap:/etc/postfix/aliases.cf to the alias_maps directive in main.cf and creating the /etc/postfix/aliases.cf file that specifies how to connect to LDAP and where the aliases are in LDAP. We restart slapd and then Postfix; we're now ready to add a mail alias. We create an LDIF file called alias.ldif and add it to the directory. That's it!

Configuring the Cyrus IMAP MDA

We'll be using the Cyrus IMAP mail delivery agent (MDA) v2.2.10. Complete configuration of the Cyrus IMAP server is beyond the scope of this article, but example working configuration files are provided in the Resources. The Cyrus IMAP server is developed by the same group who developed Cyrus SASL, so SASL and single sign-on support work as expected.

Like Postfix, Cyrus IMAP has two configuration files: /etc/imapd.conf and /etc/cyrus.conf. We'll be dealing only with /etc/imapd.conf. Again there are a few prerequisites: SSL certificate/key pair, host principal and service principal. The service principal should be called imap/mail.example.com@CI.UCHICAGO.EDU and stored in /etc/imap.keytab. To enable SSL, we define tls_ca_path, tls_cert_file and tls_key_file options, accordingly. To use SASL, we define sasl_pwcheck_method, sasl_mech_list and sasl_minimum_layer options. The values for these options are identical to those set in /etc/sasl2/smtpd.conf for Postfix. Like Postfix, Cyrus IMAP needs to be told where its keytab file is. We do this by editing /etc/conf.d/cyrus or making sure the KRB5_KTNAME environment variable is set in the init script prior to starting the IMAP dæmon. Once all this has been done, we should make sure saslauthd is running and then start the imap init script.

Wrapping Up

We certainly have covered a whole lot in a short time, but all this hard work has given you a secure and scalable corporate directory. We've just implemented a system that works for tens of users and hosts at one location all the way up to thousands spread all over the world. In my next article, we'll tackle tying Linux and Apple OS X clients into our system to see the fruits of our labor.

Acknowledgements

This work was supported by the Mathematical, Information and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract W-31-109-ENG-38. Additional support has been provided by the Computation Institute at the University of Chicago and the National Science Foundation under Grant SCI: 0451491.

Resources for this article: www.linuxjournal.com/article/8581.

Ti Leggett (leggett@mcs.anl.gov) is a systems administrator for the Futures Laboratory of the Mathematics and Computer Science Division at Argonne National Laboratory. He also has a joint appointment with the Computation Institute at the University of Chicago.

64-bit STAR-CD Compiled With PGI

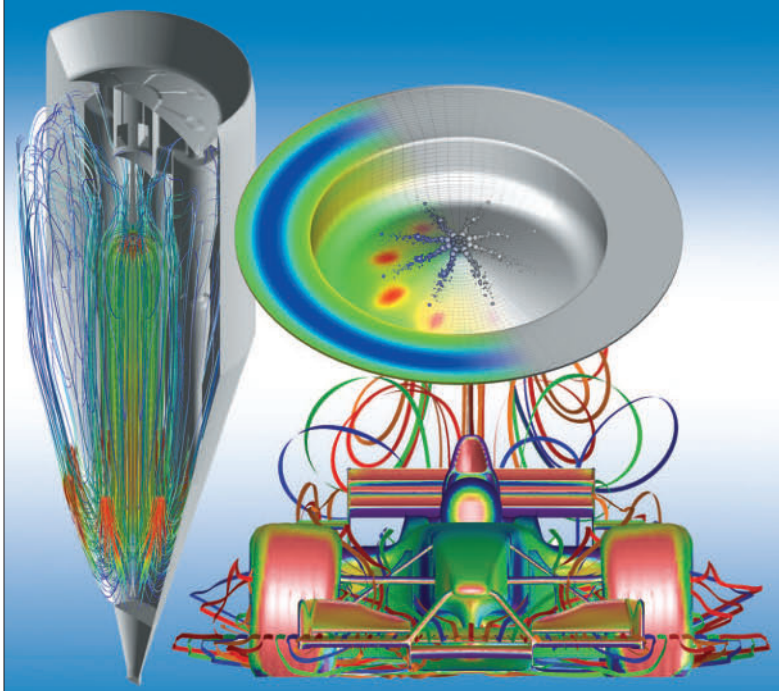

STAR-CD delivers leading CFD technology with a choice of STAR-Design or CAD-embedded options, enabling engineers to build models easily with their own choice of plug-ins. These models can be imported into the pro-STAR GUI for advanced CFD analysis using the full capabilities of STAR solvers. STAR-CD's multi-level, full spectrum approach enables enterprise-wide CFD strategies and encourages collaboration between design, application and R&D engineers.
Visit ***www.cd-adapco.com*** to learn more about STAR-CD simulation solutions for 64-bit systems.

STAR-CD software is built for AMD Opteron and Intel EM64T processor-based systems using ***PGI Compilers and Tools.***

CD-adapco

Memory Management Approach for Swapless Embedded Systems

This article presents a strategy for managing memory allocation in swapless, embedded systems to help you avoid system slowness and the dreaded Out-of-Memory killer exception.

BY MAURICIO LIN, VILLE MEDEIROS, RAONI NOVELLINO, ILIAS BIRIS AND EDJARD MOTA

The Linux kernel Out-of-Memory (OOM) killer is not usually invoked on desktop and server computers, because those environments contain sufficient resident memory and swap space, making the OOM condition a rare event. However, swapless embedded systems typically have little main memory and no swap space. In such systems, there is usually no need to allocate a big memory space; nevertheless, even relatively small allocations may eventually trigger the OOM killer.

Experiments with end-user desktop applications show that when a system has low memory—that is, it is about to reach the OOM condition—applications could become nonresponsive due to system slowness. System performance is affected when physical memory is about to reach the OOM condition or is fully occupied. System slowness should be prevented as such behaviour brings discomfort to end users.

Furthermore, the process selection algorithm used by the kernel-based OOM killer was designed for desktop and server computer needs. Thus, it may not work properly on swapless embedded systems, because at any moment it can kill applications that a user may be interacting with.

In this article, we present an approach that employs two memory management mechanisms for swapless embedded systems. The first is applied to prevent system slowness and OOM killer activation, by refusing memory allocations based on a predefined memory consumption threshold. Such a threshold should be determined and calibrated carefully in order to optimize memory usage while avoiding large memory consumption that may lead to system delay and invocation of the OOM killer. We call it the Memory Allocation Threshold (MAT).

The second mechanism employs an additional threshold value known as the Signal Threshold (ST). When this threshold is reached, the kernel sends a low memory signal (LMS), which should be caught by user space, triggering memory release before crossing the MAT. Both thresholds are implemented by a kernel module, the Low Memory Watermark (LMW) module. We offer some experimental results that point out situations when our approach can prove useful in optimizing memory consumption for a class of embedded systems.

Memory Management Approach

System performance is degraded when the memory requirements of active applications exceed the physical memory available on a system. Under such conditions, the perceived system response can be significantly slow. On swapless devices, application memory needs can drive the system to such conditions often, because system internal main memory is low and the chance of applications occupying the whole physical memory is high.

Memory resources should be managed differently on such devices to avoid slow system responsiveness. The memory allocation failure mechanism can be applied to prevent slowness. Preventing system slowness makes OOM killer invocation rare. Thus, such a mechanism also can reduce the chances of triggering the OOM killer, whose process selection algorithm may choose an unexpected application to be killed on devices with low memory and no swap space.

Memory allocation failure means refusing memory allocations requested by applications. It is carried out according to a MAT value that is set based on experimentation with various use cases of end-user applications. MAT should be set sufficiently high to allow applications to allocate necessary memory without affecting overall system performance, but its value should be well defined to guarantee memory allocation failure when necessary to prevent extreme memory consumption.

Before memory allocation failure occurs, however, process termination can be performed to release allocated memory. It can be triggered by transmitting the LMS from kernel space to user space to notify applications to free up memory. LMS is

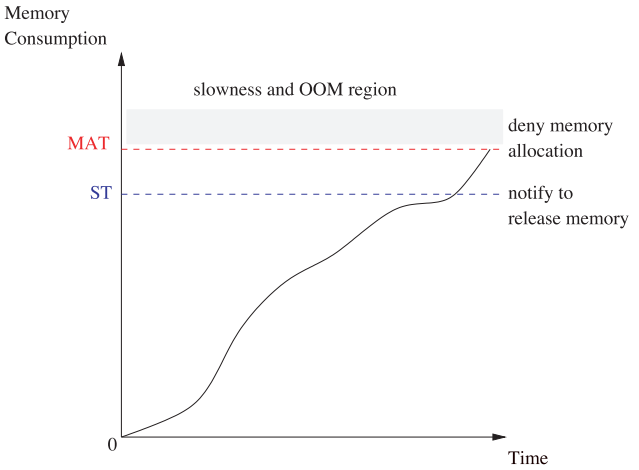

Figure 1. Signal Threshold is smaller than Memory Allocation Threshold.

Find out why some of the world's largest internet sites choose Open Source Storage for their enterprise data center needs.

Servers
Stability
Storage
Support
Savings
Speed

The VS1100 server includes the following items below:
2x 246 AMD 1MB 32/64-bit Opteron CPU's
S2881 Motherboard with:
- **On-board Video**
- **Dual LAN 10/100/1000**
- **Supports SATA RAID 0/1**
- **Supports Dual Core Opteron CPU's**

2x 512MB PC3200 ECC REG Memory (1GB total)
1x 80GB SATA Hard Drive
2x Hot-Swap Drive Bays
1U Black rackmount case with 350 Watt Power Supply
3 Year Parts and Labor Warranty
30 Day Money Back Guarantee
$1299.99 (Limited to Stock On Hand).

Specializing in high performance computing and complete data center implementation

1-866-664-STOR
1900 Wyatt Dr. Bldg.16&17
Santa Clara, CA 95054
http://www.OSShpc.com

Open Source Storage

©2002-2005 Open Source Storage, Inc. All Rights Reserved. AMD, AMD Opteron, combinations thereof, are trademarks of Advanced Micro Devices, Inc.

dispatched according to ST value. ST should be smaller than MAT, as shown in Figure 1, because the LMS should occur well before memory allocation failure.

If the LMS dispatch is successful and memory is released by receiving the signal, a possible memory allocation failure will be prevented. A useful scenario could involve running some window-based applications, A, B and C, consuming chunks of memory, while their window frames can superimpose one another (assuming the use of a simple window manager environment such as Matchbox). Assuming that application A is the one the user is interacting with at the moment MAT is reached, instead of denying memory allocation to A, it would be preferable to attempt to free up memory allocated by applications B and C, which are not visible to the user. Doing this would allow the user to continue working with application A.

However, memory allocation failure could be unavoidable for some application use cases. For instance, such a case could involve a single window-based application, consuming memory at a constant rate, that the user is interacting with. Releasing memory from other applications would not be as desirable in this situation, because there may be no other window-based applications from which memory could be released. Therefore, a more desirable solution would be to fail memory allocation requested by the guilty application, selecting it as a candidate for termination.

In our proposal, the kernel should provide two mechanisms to deal with management of memory in extreme cases of low memory levels:

- Failure of brk(), mmap() and fork() system calls: deny memory allocation requests to prevent system slowness and kernel OOM killer invocation according to a previously calibrated MAT level.

- Low memory signal: Kernel Event Layer signal sent by the kernel to a user-space process terminator, which should employ a process selection algorithm that works based on a specified ST.

Using these mechanisms, it would be possible to identify when memory can be released or when to deny further allocations. Denying memory allocations should happen only when memory release attempts cannot be successful.

Low Memory Watermark (LMW) Module

LMW is a kernel module based on the Linux Security Module (LSM) framework. It implements a heuristic to check the physical memory consumption threshold for denying memory allocation and notifying user space to free up memory. A user-space process terminator can be employed to free up memory. Formulas for low memory watermark thresholds are as follows:

- deny_threshold = physical_memory * deny_percentage

- notify_threshold = physical_memory * notify_percentage

physical_memory is the system's main memory and is represented by the kernel global variable totalram_pages. deny_percentage and notify_percentage are tunable kernel parameters, and the value of these can be altered through the sysctl interface. These parameters are bound to the /proc filesystem and can be written to and read from, using standard commands such as echo and cat. These parameters may be handled as follows:

```
$ echo 110 > /proc/sys/vm/lowmem_deny_watermark
```

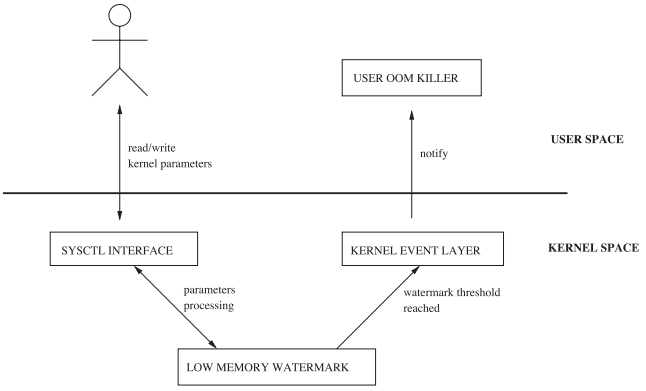

Figure 2. Low Memory Watermark Architecture

Listing 1. Algorithm of MAT and ST Watermarks Heuristic

```
1 static int low_vm_enough_memory(long pages)
2 {
3   unsigned long committed;
4   unsigned long deny_threshold, notify_threshold;
5   int cap_sys_admin = 0;
6
7   if (cap_capable(current, CAP_SYS_ADMIN) == 0)
8     cap_sys_admin = 1;
9
10  if (deny_percentage==0||notify_percentage==0)
11    return __vm_enough_memory(pages,cap_sys_admin);
12
13  deny_threshold=
14        totalram_pages*deny_percentage/100;
15  notify_threshold=
16        totalram_pages*notify_percentage/ 100;
17
18      vm_acct_memory(pages);
19      committed = atomic_read(&vm_committed_space);
20      if (committed >= deny_threshold) {
21              enter_watermark_state(1);
22              if (cap_sys_admin)
23                      return 0;
24              vm_unacct_memory(pages);
25              return -ENOMEM;
26      } else if (committed >= notify_threshold) {
27              enter_watermark_state(1);
28              return 0;
29      }
30      enter_watermark_state(0);
31      return 0;
32 }
```

Is your data center cramping your style?

Growth always seems like a good idea. An extra processor here—one more server there. Until, all the sudden your data center feels as crowded as a center seat in coach. Let the Penguin upgrade you. Penguin Computing introduces BladeRunner™ 4140 the industry's densest Linux blade server. It comes with the AMD Opteron™ HE processor, which offers simultaneous 32- and 64- bit computing. So now you can pack 48 cores into a miniscule 4U of rack space, and optimize your data center. And put that 8GB of PC3200 RAM per blade to work and run your 64-bit apps in a fraction of the space. So go ahead. Stretch your legs. Tilt your seat back. Love what you do.

Visit www.penguincomputing.com

AMD 64 Opteron

AMD Opteron is a trademark of Advanced Micro Devices, Inc
Other names are for information purposes only and may be trademarks of their respective owners.

Penguin Computing is a registered trademark of Penguin Computing, Inc. Linux is a registered trademark of Linus Torvalds.

```
$ echo 90 > /proc/sys/vm/lowmem_notify_watermark

$ cat /proc/sys/vm/lowmem_deny_watermark
110

$ cat /proc/sys/vm/lowmem_notify_watermark
90
```

The LWM architecture is illustrated in Figure 2. Basically, LWM overrides the kernel default overcommit behaviour by setting the vm_enough_memory function pointer field in the security_operations structure to point to the function low_vm_enough_memory(). low_vm_enough_memory() implements a heuristic based on the formula described earlier. Binding vm_enough_memory to low_vm_enough_memory() permits interception of all requests for allocation of memory pages in order to verify whether the committed virtual memory has reached the MAT or ST watermarks. Listing 1 presents how the MAT and ST watermarks are implemented in the low_vm_enough_memory() function.

The code in Listing 1 is explained below:

- Lines 7, 8: verify whether the current process has root privileges.
- Lines 10, 11: if MAT or ST watermarks are zero, perform the default overcommit behaviour.
- Lines 13–16: calculate the low memory watermark thresholds.
- Line 18: the pages are committed to update the amount of vm_committed_space.
- Line 19: the amount of committed memory is acquired.
- Line 20: verify whether committed memory has reached the MAT watermark.
- Line 21: set a flag state to 1 if MAT has been reached—state=1 means any (or both) of the two thresholds was reached.
- Lines 22, 23: do not deny memory allocation for root programs—allocation is successful for these.
- Line 24: uncommit the current committed pages since MAT was reached.
- Line 25: return no memory available message.
- Line 26: verify whether committed memory has reached the ST watermark.
- Lines 27, 28: set the state to 1, and allocation has succeeded.
- Line 30: set the state to 0 (if no threshold was reached).
- Line 31: memory allocation has succeeded.

The enter_watermark_state() function determines whether the low memory watermark condition has been reached and eventually sends the LMS to user space. A global boolean variable, lowmem_watermark_reached, marks the state of entering or exiting from low memory watermark conditions, being assigned to a value of 1 or 0, respectively. LMS is dispatched whenever a change in the value of this variable occurs.

Listing 2. Algorithm of Entering Watermark States

```
1 static void enter_watermark_state(int new_state)
2 {
3   int changed = 0, r;
4
5   spin_lock(&lowmem_lock);
6   if (lowmem_watermark_reached != new_state) {
7     lowmem_watermark_reached = new_state;
8     changed = 1;
9   }
10  spin_unlock(&lowmem_lock);
11  if (changed) {
12    printk(KERN_DEBUG MY_NAME ": changed to %d\n",
13           new_state);
14    r = kobject_uevent(&kernel_subsys.kset.kobj,
15                       KOBJ_CHANGE,
16                       &low_watermark_attr.attr);
17    if (r < 0)
18      printk(KERN_ERR MY_NAME
19             ": kobject_uevent failed: %d\n", r);
19    }
20  }
```

Listing 2 illustrates how the state is changed, and the LMS is sent to user space. Intuitively, the code works as follows:

- Line 5: lock to avoid a race condition.
- Line 6: verify whether the new state is different from the old one.
- Lines 7, 8: update the lowmem_watermark_reached and changed variable.
- Line 10: unlock to leave the critical region.
- Line 11: verify whether the state was changed.
- Lines 12–16: log that the state was modified and send the signal using the Kernel Event Layer mechanism.
- Lines 17–19: log a message if an error occurred.

Tuning Memory Consumption Parameters

Tuning MAT can be done empirically based on some use cases. Tuning of the ST watermark is not presented here, but it is usually done in the same manner as MAT. Applications used in the scenarios involved should succeed in filling the memory totally, thus overloading the system. Doing this can trigger system slowness and kernel OOM killing, thus ensuring a valid use

case for tuning the MAT watermark.

As discussed previously, an optimal MAT value, the memory allocation refusal threshold, should be such so as to avoid system slowness and kernel OOM killer execution. MAT value is given in terms of the percentage of memory that the kernel commits, possibly reaching values more than 100% due to Linux kernel's memory overcommit feature.

Basically, three behaviours need to be identified during experimentation: OOM killer execution, refusal of memory allocation and system slowness. The experiments were performed using a swapless device with 64MB of RAM memory and 128MB of Flash memory. The Flash memory is the secondary storage used as a block device to retain data.

The first use case involves reaching the MAT in a gradual manner, running the following applications (in the order they are listed): Web browser, e-mail client, control panel to configure the system and image viewer. First, the Web browser loads a Web page, followed by the e-mail client loading some 360 messages in the inbox, followed by the control panel, which is simply opened, and finally by the image viewer loading a number of image files, one after the other (only one image is loaded to memory at a time). Each image file is progressively larger than the previous one, all a few hundreds of KB, but one is about 2MB. Loading these files progressively can cause a different system behaviour according to different MAT values. Table 1 illustrates the results of this scenario when varying the MAT values.

Table 1. MAT Value for Web Browser, E-mail Client, Control Panel and Image Viewer Use Case

MAT (%)	OOM Killer	Denied Memory	Slowness
120	2	0	3
119	r	0	1
115	5	0	0
112	2	7	1
111	0	5	0
110	0	5	0

A MAT threshold of 120% is not a good choice, because it allows OOM killing to occur twice while slowness occurs three times. The best MAT value, in this use case, is 111%, because at that level the system is able to deny all memory allocations preventing system slowness and kernel OOM killer execution.

In the use case described above, whenever the OOM killer occurs, it always kills the image viewer application. Slowness takes place when the image viewer tries to load the heavy image file of 2MB. During the experiment, it was perceived that the OOM killer is always started during the system slowness, and usually system slowness is so severe that waiting for OOM killing is not viable.

A second use case could try to reach the MAT threshold in a more direct manner. The following applications are started:

Web browser, PDF viewer, image viewer and control panel. The Web browser loads a Web page, then the PDF viewer attempts to load a file of 8MB, followed by the image viewer loading an image file of 3MB and finally invoking the control panel.

In this use case, whenever the image viewer loads the image file, the PDF file of 8MB loaded previously is unloaded, because of the ST threshold being reached, causing a signal dispatch to user space in order to free up memory. The observed behaviour also involved the termination of the control panel application, which can be attributed to memory allocation denial due to having reached MAT. Table 2 presents the experimental results for this use case for different MAT values.

Table 2. MAT Value for Web Browser, PDF Viewer, Image Viewer and Control Panel Use Case

MAT (%)	OOM Killer	Denied Memory	Slowness
120	0	0	5
113	0	0	5
112	0	1	4
111	0	2	3
110	0	5	0

This use-case scenario indicates a reliable MAT value of 110%. Slowness occurs for values above 110% when the control panel is started. Figure 3 illustrates how the MAT and ST behave in this use case. The memory consumption curve shown is assumed, but it does not in any way alter the aforementioned results.

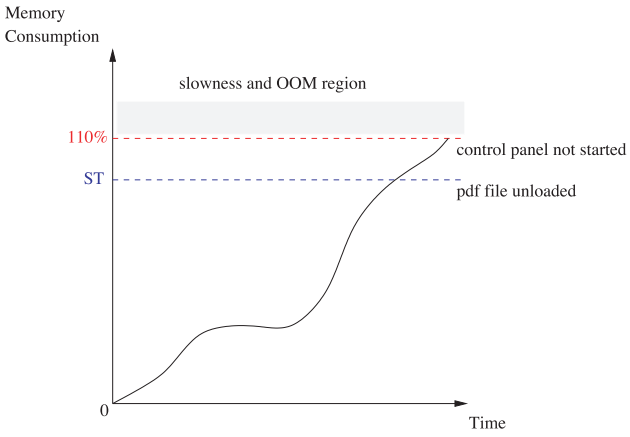

Figure 3. Low memory watermark graphic, based on Web browser, PDF viewer, image viewer and control panel use case.

During experimentation, it is important to verify whether the planned use cases are satisfactory for calibrating the MAT value, because there could be use cases that do not overload memory allocations. An example of such a scenario could be invoking the Web browser to download a file of 36MB in the background while playing a game at the same time. Our experiments indicated that this use case was not as useful in determining a realistic MAT value, because it worked successfully even with a MAT value of 120% or higher.

Some Additional Remarks

A useful approach in assisting the fast selection of processes to be killed, in order to release memory, could involve registering applications as killable or not. Applications considered killable could be registered on a list known as the Red List. Additionally, other applications, crucial for the correct functionality of the system, such as the X Window System, should not be killed under any circumstances and could be registered on a list known as the White List.

End users could be allowed to choose which applications should be registered on the Red or White Lists. However, this would require a security mechanism in place to ensure that applications on the Red List or White List do not cause any unexpected conditions or instabilities. If application A is the culprit by consuming tons of memory continuously, it cannot be on the White List. Likewise, if killing application B can break down overall system functionality, then it cannot be on the Red List. A heuristic could be employed for selecting in advance which applications can be registered on the Red List or White List. Preselected applications could then be presented to the user to be opted for registration on the respective list, thus improving user-friendliness while avoiding potential problems from choosing erratically.

The Red List and White List could be implemented in kernel space, with each list also reflected in the /proc filesystem. ST can be used to notify user space the moment when the Red and White Lists should be updated. Afterward, the kernel can start terminating applications registered on the Red List in order to release memory. Perhaps a ranking heuristic can be employed in kernel space to prioritise entries on the Red List. Figure 4 illustrates a possible architecture of OOM killer, based on Red List and White List approach. If it is not enough simply to kill processes on the Red List, other processes, not appearing on the White List, could be killed as well, as a last measure to ensure system stability.

It is interesting to maintain a mechanism based on having one heuristic for selection and termination of processes in user space and another one in kernel space, because each space can offer different pieces of information that may prove useful to the ranking criteria. For instance, in user space it is possible at any moment to know which window-based applications are active, that is, visible and used by the end user, but in kernel space such information is not as easily attainable. Hence, if there is a heuristic that needs to verify whether any window-based application is active or not, it should be implemented in the user space.

Conclusion

Dealing with swapless embedded systems requires establishing an alternative memory management approach, in order to prevent slowness and to control OOM killer invocation and execution. The idea based on MAT and ST is simple yet practical and tunable on different swapless embedded devices, because the LMW kernel module provides the /proc and sysctl interfaces to change the MAT and ST values from user space as necessary.

Additional mechanisms can be implemented, such as the Red

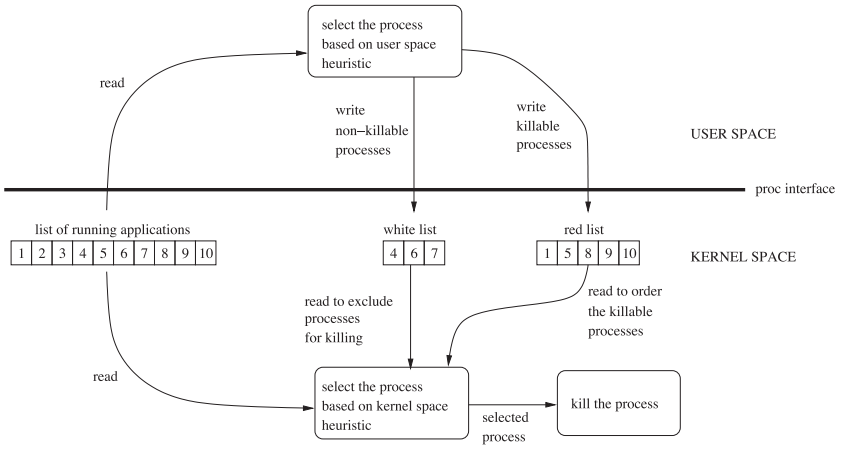

Figure 4. Architecture of OOM Killer Based on the Red List and White List Approach

and White registration Lists. It is also interesting to design different selection criteria that take into account features related to swapless embedded devices.

Acknowledgements

We wish to offer our thanks to Kimmo Hämäläinen of Nokia Research Center (NRC) and Monica Nogueira for help on the organization and contents of this document, Juha Yrjölä of NRC for help with Low Memory Watermark coding and Fabritius Sampsa of NRC for providing us with the opportunity to develop this work.

Mauricio Lin is a software engineer at Nokia's Institute of Technology (INdT) located in Manaus/Brazil. Mauricio has worked with Linux memory management for embedded systems since 2003. He also has contributed to the proc-pid-smaps.patch of the Linux kernel related to memory consumption analysing. Mauricio received a BSc in Data Processing from the Federal University of Amazonas. He has practiced Kung Fu - Wushu since childhood and has graduated in the Northern Shaolin style. He can be reached at mauriciolin@gmail.com.

Ville C. L. de Medeiros has been passionate about Linux since 1997. He started as a trainee administrator at the Department of Computer Science, Federal University of Amazonas (UFAM) and became the university's network manager after graduation. He then took the initiative to convert all network services from mainframe-based to Linux. He currently works in the Linux Embedded Lab in Nokia's Technology Institute at Manaus/Brazil. He can be reached at ville.medeiros@gmail.com.

Raoni Novellino has been a Linux user and developer for two years and is currently working in Nokia's Technology Institute at Manaus/Brazil. He can be reached at rnovellino@gmail.com.

Ilias Biris holds a PhD in Artificial Intelligence from the University of Edinburgh and has been a Linux user and developer for ten years. He works as project coordinator in the Linux Embedded Lab of Nokia's Institute of Technology in Manaus/Brazil and can be reached at xyz.biris@gmail.com. Apart from Linux, Ilias enjoys early morning Tai Chi and a good duel in Kendo.

Edjard Mota holds a PhD in Artificial Intelligence from the University of Edinburgh and has been an avid Linux user for ten years and a developer for two. He is currently managing the Linux Embedded Lab of Nokia's Technology Institute at Manaus/Brazil. He is an early-bird and enjoys doing Yoga to start the day. He can be reached at edjard.mota@gmail.com.

$119 qty 100
- 200 MHz ARM9
- 10/100 Ethernet
- PC/104 bus

TS-7200 ARM9 Single Board Computer

Linux 2.4

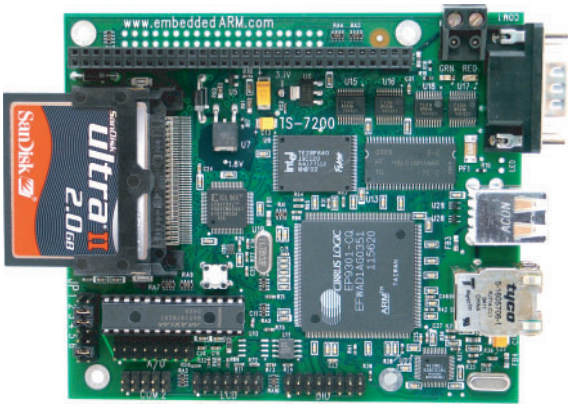

Shown with optional Compact Flash

$149 qty 1
- 32 MB SDRAM (64 MB optional)
- 8 MB Flash (16 MB optional)
- Compact Flash
- 10/100 Ethernet
- 2 USB ports
- 20 Digital I/O
- 2 Serial Ports

Options:
- RS-485
- 8 ch 12-bit A/D
- RTC (battery-backed)

- Boots Debian stable from Compact Flash
- Boots TS-Linux from on-board Flash

Technologic SYSTEMS

We use our stuff.
Visit our TS-7200 powered website at

- Many x86 and ARM based SBCs and peripherals available
- Call for custom designs

(480) 837-5200
www.embeddedARM.com

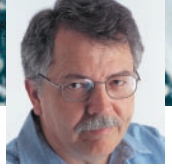

The World Live Web

Doc Searls continues to bring us leading-edge ideas from the forefront of the Web. **BY DOC SEARLS**

There's a split in the Web. It's been there from the beginning, like an elm grown from a seed that carried the promise of a trunk that forks twenty feet up toward the sky.

The main trunk is the static Web. We understand and describe the static Web in terms of real estate. It has "sites" with "addresses" and "locations" in "domains" we "develop" with the help of "architects", "designers" and "builders". Like homes and office buildings, our sites have "visitors" unless, of course, they are "under construction".

One layer down, we describe the Net in terms of shipping. "Transport" protocols govern the "routing" of "packets" between end points where unpacked data resides in "storage". Back when we still spoke of the Net as an "information highway", we used "information" to label the goods we stored on our hard drives and Web sites. Today "information" has become passé. Instead we call it "content".

Publishers, broadcasters and educators are now all in the business of "delivering content". Many Web sites are now organized by "content management systems".

The word content connotes substance. It's a material that can be made, shaped, bought, sold, shipped, stored and combined with other material. "Content" is less human than "information" and less technical than "data", and more handy than either. Like "solution" or the blank tiles in *Scrabble*, you can use it anywhere, though it adds no other value.

I've often written about the problems that arise when we reduce human expression to cargo, but that's not where I'm going this time. Instead I'm making the simple point that large portions of the Web are either static or conveniently understood in static terms that reduce everything within it to a form that is easily managed, easily searched, easily understood: *sites*, *transport*, *content*.

The static Web hasn't changed much since the first browsers and search engines showed up. Yes, the "content" we make and ship is far more varied and complex than the "pages" we "authored" in 1996, when we were still guided by Tim Berners-Lee's original vision of the Web: a world of documents connected by hyperlinks. But the way we value hyperlinks hasn't changed much at all. In fact, it was Sergey Brin's and Larry Page's insights about the *meaning* of links that led them to build Google: a search engine that finds what we want by giving maximal weighting to sites with the most inbound links from other sites that have the most inbound links. Although Google's PageRank algorithm now includes many dozens of variables, its founding insight has proven extremely valid and durable. Links have value. More than anything else, this accounts for the success of Google and the search engines modeled on it.

Among the unchanging characteristics of the static Web is its nature as a haystack. The Web does have a rudimentary directory with the Domain Name Service (DNS), but beyond that, everything to the right of the first single slash is a big "whatever". UNIX paths (/whatever/whatever/whatever/) make order a local option of each domain. Of all the ways there are to organize things—chronologically, alphabetically, categorically, spatially, geographically, numerically—none prevails in the static Web. Organization is left entirely up to whoever manages the content inside a domain. Outside those domains, the sum is a chaotic mass beyond human (and perhaps even machine) comprehension.

Although the Web isn't organized, it can be searched as it is in the countless conditional hierarchies implied by links. These hierarchies, most of them small, are what allow search engines to find needles in the World Wide Haystack. In fact, search engines do this so well that we hardly pause to contemplate the casually miraculous nature of what they do. I assume that when I look up linux journal diy-it (no boolean operators, no quotes, no tricks, just those three words), any of the big search engines will lead me to the columns I wrote on that subject for the January and February 2004 issues of *Linux Journal*. In fact, they probably do a better job of finding old editorial than our own internal searchware. "You can look it up on Google" is the most common excuse for not providing a search facility for a domain's own haystack.

I bring this up because one effect of the search engines' success has been to concretize our understanding of the Web as a static kind of place, not unlike a public library. The fact that the static Web's library lacks anything resembling a card catalog doesn't matter a bit. The search engines are virtual librarians who take your order and retrieve documents from the stacks in less time than it takes your browser to load the next page.

In the midst of that library, however, there are forms of activity that are too new, too volatile, too unpredictable for conventional Web search to understand fully. These compose the *live Web* that's now branching off the static one.

The live Web is defined by standards and practices that were nowhere in sight when Tim Berners-Lee was thinking up the Web, when the "browser war" broke out between Netscape and Microsoft, or even when Google began its march toward Web search domination. The standards include XML, RSS, OPML and a growing pile of others, most of which are coming from small and independent developers, rather than from big companies. The practices are blogging and syndication. Lately podcasting (with OPML-organized directories) has come into the mix as well.

These standards and practices are about time and people, rather than about sites and content. Of course blogs still look like sites and content to the static Web search engines, but to see blogs in static terms is to miss something fundamentally different about them: they are alive. Their live nature, and their humanity, defines the *live* Web.

It is essential that we understand the live Web on its own terms, rather than those leveraged from the static Web.

Blogs are journals, not sites. They are written, not built. The best ones have a heart that beats daily or faster. The writing itself is more conversational than homiletic (which is how I'm behaving here, in a print publication with a monthly heartbeat). That means its authors are *speaking*, and not just "creating content". They speak to readers and other bloggers who speak back, through e-mails, comments or on blogs of their own. That means what each blogger says is often incomplete and provisional. Like all forms of life, blogging remains unfinished for the duration. (Site content, on the other hand, is finished at any one time, then replaced with other finished content.)

A few months back, I was asked to explain blogging to somebody who knew nothing about it. When I finished, the guy understood that blogging was a new form of journalism that gave individuals a higher degree of leverage than ever before. He then instructed me, as a fairly well-known blogger, to devote my remaining life immediately to correcting the familiar evils of the world.

I replied that I was already 57 years old and tired of pushing large rocks up steep hills for short distances—also of getting flattened by the rocks that rolled back over me. I told him blogging might make Sisyphus' life a bit easier in some cases, but that its better leverage was on snowballs. My work as a blogger, I explained, is rolling snowballs downhill. Some I create new; others I push along, adding a small measure of mass along the way.

My point: rolling snowballs is way different from building sites and transporting content. Not totally different, perhaps, but enough to fork the Web.

Blogging predated syndication, but it was syndication that began to give form to the live Web. Syndication provided a way for people, and the tools they use, to pay attention (through subscription) to feeds from syndicated sources. At first these sources were blogs and publications, but later they came to include searches for topics of conversation, including the names of authors, URLs and permalinks for particular blog posts or news stories. Many of those sources were not the blogs themselves, but search engines reporting the results of keyword and URL searches.

At the time of this writing, the most popular live Web search engine is Technorati (now about #700 on Alexa, with around 80-million page views per day). It was born in November 2002 on a Linux box from Penguin Computing that sat in David Sifry's basement. The box was loaned to help the two of us write a feature on blogging that ended up running in the February 2003 issue of *Linux Journal*. David wrote Technorati to help him do research for the story. The first time I saw it, I also saw the fork in the Web. What Technorati searched was alive, moving, changing. Its results were also radically different from what I got from the static Web. This past spring somebody who works for Victoria's Secret complained to a friend about the limited knowledge the company had obtained regarding its IPEX bra,

ONCE AGAIN,
HEAP PROBLEMS HAD
SPOILED CODY'S DAY

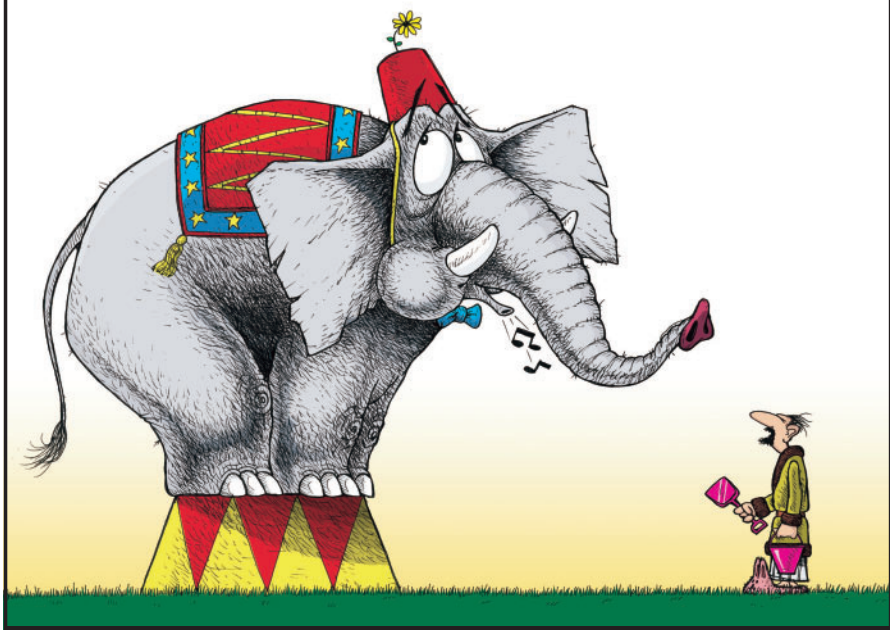

Characters and Images ©2004 Brad Fitzpatrick, ActiveEdge. All Rights Reserved.

Debugging heap allocation problems can be a real chore, but TotalView now has built-in memory features that track memory usage for all processes and can even stop execution at the point that a memory problem occurs. And it's all integrated, so there's no need to interrupt your debug session to invoke an external memory tool. Etnus TotalView is also the best threads debugger available and offers superior C++ support. So, don't forget to download a free fully functional trial of TotalView today.

Try TotalView FREE at www.etnus.com

TotalView, the Most Advanced Debugger on Linux and UNIX

NOW ON Mac OS X

ETNUS TOTALVIEW

which had hit the stores only a few weeks before. A search on Google brought up only Victoria's Secret's own site and a few others that offered retail information. My friend showed her a Technorati search for "ipex" that brought up hundreds of posts, mostly by women telling to other women how much they liked the bra. That search was a window on Unfiltered Truth that barely resembled anything the company would get from focus groups or other customary forms of market research.

Today there are a half-dozen engines devoted to searching the live Web. They're all different. Blogpulse stresses trending and ranking (with a great UI and excellent graphics). PubSub doesn't offer Web search but instead concentrates on keyword search feeds to users' aggregators. Bloglines integrates search with aggregation and other services. IceRocket emphasizes performance and simplicity. Technorati focuses on rapid indexing, tag search and hot topics. Feedster leads with personalization and index size.

All those characterizations are simplistic and incomplete. They are also obsolete by the time you read this. The whole category is changing as rapidly as the individuals and social trends they follow, as well as the technologies that make them possible and the developers who do new things with those technologies. A couple days ago I talked with a new company that gathers and syndicates conversation around local businesses and services, making the Live branch of the Wide Web as local as possible. I have at least one of these conversations every week.

This morning I had a conversation with some techies involved in "microformats". These are described on the microformats.org site as "a set of simple, open data formats built upon existing and widely adopted standards. Instead of throwing away what works today, microformats intend to solve simpler problems, first by adapting to current behaviors and usage patterns (for example, XHTML, blogging)." Rather than specifications and standards, microformats are "design principles", "methods of adaptation to usage patterns", "correlated with semantic XHTML and the Real World" and "a way of thinking about data". Far as I know, nobody around microformats wants to patent them or to patent a business model that makes use of them. Just as nobody patented RSS (which first meant "rich site summary" but came to mean "really simple syndication" after Dave Winer led its evolution into a stable live Web enabler). We can thank this kind of largesse for the Net and the Web, as well as for Linux and the Free Software and Open Source movements.

Tagging is a perfect example of standards and practices evolving in a live, organic way. Tags are labels that serve as categories, attached by users to photographs, lists, blog posts or anything they put up on the Web (or that others put up). Tags first appeared on del.icio.us, a social bookmarks manager, and on Flickr, a photo sharing service. In both cases, developers put users in control of their own creations (note that I avoid saying "content") and the descriptions of those creations. Later, Technorati began doing not only tag searches, but also establishing standards for tagging in links (including the rel="tag" element). Authors and users began adding tags to all kinds of stuff. As a result, tags are now becoming a form of live Web organization.

The blogging branch of the live Web has another kind of order: chronological. Whether served up by TypePad or Drupal or Manila or some other system, blogs are all organized the same way: *blogname.suffix/year/month/day/post*. The permalink of the post is its unique URL.

Any pile of organized data can be archived. This means that the part of the Web that's least static is also the part that can be archived and organized like a library—and researched the same way, only better. Think about the amount of data that can be gathered from a sum of sources organized by date and category (tags). Think of the intelligence that can be gleaned from that. Also think about the business there might be in facilitating or selling that intelligence.

I see by Netcraft that all the live Web search engines I've named so far run on Linux. So do Google, AskJeeves and A9. Even MSN Search runs on Linux, through Akamai's giant server farms. The only exception is Yahoo, running its own breed of BSD (which is still an open-source OS).

As I write this, I'm also helping put together the Syndicate conference in San Francisco (December 12–14, 2005, at the Hilton downtown—this issue of *Linux Journal* should be on the newsstands at that time). It is customary at tradeshows to look to vendors and large service providers for leadership. With the live Web, however, leadership doesn't just come from the big guys. In fact, most of it comes from independent developers and pioneering users. In this respect, the live Web is more an ecosystem than an industrial category. The folks standing on stage will have lots to say, but so will the folks who compose what we used to call "the audience". It will be interesting to see how conversations go.

It also will be interesting to see which way the live Web carries Linux innovations and conversations about them. Linux and open-source development have always had their live qualities. As the live Web grows, we can expect those to become more organized (by chronology or tag, for example) at the very least.

Is it possible that "live" will join "free" and "open" in our pantheon of adjectives? Possibly. Whether or not it does, I'd like to thank my son Allen for being the first to utter "World Live Web", providing me with a perspective I never knew I lacked, until I heard it.

His original vision of the World Live Web was a literal one: a Web where anybody could contact anybody else and ask or answer a question in real time. When he first encountered the Web, as a researcher, he saw it as something fundamentally deficient at supporting the most human forms of interaction: the kind where one person increased the knowledge of another directly.

We've moved a long way in the live direction since Allen first introduced me to the concept. VoIP alone is a huge live category. Mobile Web progress will all happen along its live branch.

Where it goes exactly is anybody's guess. All we can say for sure is it's headed toward the sky.

Doc Searls is Senior Editor of *Linux Journal*.

Polywell High Performance Systems

Reliable 64-bit Computing with First Class Service

AMD Dual-Core technology Enables one platform to meet the needs of multi-tasking and multi-threaded environments; providing platform longevity

Allows users to run 32- and 64-bit applications as they desire – without sacrificing performance

1U 4-way, 64GB DDR, 2TB RAID

- 2 x AMD® Opteron™ Dual-Core Processors 275+
- with Hyper Transport Technology
- Upto 64GB 400MHz ECC DDR (16 Sockets)
- Upto 2TB 4 x 500G Swap Drive per 1U Rack
- 1 x PCI-X 133MHz, 1 x PCI-E Solt
- Dual Gigabit Ethernet, ATI Graphics, 4x SATA-RAID
- Slim CD-ROM Drive, Optional DVD-RW or CDRW
- Optional Slim Floppy Drive
- 1U 24" Depth Rack Chassis with upto 600W PS
- 4 x Swappable Drive Bays (SATA or SCSI)
- Supports Linux, FreeBSD or Windows
- Custom Configuration Available
- Please call for other Options

1114AIS-2050M, 275+,32GB, 2TB **$8,888**
1114AIS-2050M, 265+,16GB, 1TB **$6,499**
1112ES-2200A, 242+, 2GB, 500G **$1,999**

Linux Appliance PCs

- Custom Made Odd Size Chassis
- AC or DC Power Supply
- Low Power Voltage AMD Sempron™ Processor
- or High Performance AMD Athlon™ 64 Processor
- Diskless or Flash OS Boot Drive
- Swapable Hard Drive, CD-ROM, FDD
- Integrated Graphics, Ethernet, USB
- Optional LCD LED Control Module
- IS2 Audio, MPEG2/4 Hardware Video
- Upto 4 Ethernet Ports or 4 Serial Ports
- We have over 18 years OEM Experience in Set top Box, Digital Media Player, POS Kiros, Thin Client, Networking Appliance,

SX2500SPLJ11B
OEM Appliance starts at **$299**

64GB RAM 4-way Workstation

- 2 x AMD® Opteron™ Dual-Core Processors 275+
- Upto 64GB 400MHz ECC DDR (16 Sockets)
- 2 x 133MHz PCI-X, 1 x PCI-E x16 Slots
- Dual Gigabit Ethernet, 4 x SATA-RAID Controller
- Quiet and Cool 12-Bay Tower + 600W 80Plus P/S
- 250G HD, DVD-RW, Floppy, Optional Card Reader
- 8-Layer Motherboard with Special Quiet Cooling
- On-board ATI Graphics, Sound Blaster 7.1 Sound
- Supports 64/32-bit Linux, FreeBSD or Windows
- Special for Large Memory Intensive Applications
- Built-to-Order or Configure-to-Order

64G RAM, 2x275+, QuadroFX4500 **$22,999**
32G RAM, 2x250+, QuadroFX3400 **$9,999**
16G RAM, 2x244+, QuadroFX1300 **$3,999**

1U Power Saving ISP Server

- AMD Sempron™ or Opteron™ Processor
- 512M DDR 400MHz Memory
- 80GB Hard Drive
- 10/100Mbit Ethernet
- We provide Drive Image Service
- 1U 14" Short Rack, allow 2 x 1U per Rack
- Low Power Usage to save your Data Center Cost
- Perfect Entry Level ISP Server or Appliance System
- IDE Flash Drive is available
- Supports Linux, FreeBSD or Windows
- Custom Configuration is Available
- Please call us to discuss your specification

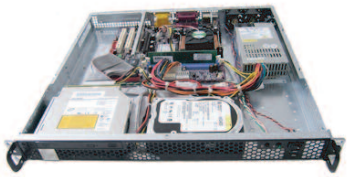

Order# VX2500SP1U-17LJ11A
starts at **$399**

2U 8-way, 5U 16-way Opteron

- 8 or 4 AMD® Opteron™ Dual-Core Processors 865+ with Hyper Transport Technology
- Upto 128GB DDR Memory for 16-way (32 sockets)
- Upto 64GB DDR Memory for 8-way (16 sockets)
- 4x Gigabit LAN, 8x SATA RAID-5 for 16-way
- 2x Gigabit LAN, 4x SATA, U320 SCSI for 8-way
- 4 x 133/100/66MHz PCI-X Slots for 16-way
- 2 x 133, 2 x 66MHz PCI-X, 1x PCI Slots for 8-way
- On-board ATI Graphics, USB 2.0
- 5U 26" Rack 1300W 3+1 Redundant P/S 16-way
- 2U 27" Rack 700W PFC P/S for 8-way
- Supports Linux, FreeBSD or Windows
- Custom Configuration Available
- Please call for other Options

5U 16-way 865+, 32GB, 2TB, 8800U5 **$29,999**
2U 8-way 865+, 4GB, 750GB, 8422B **$10,999**

SANNAS 4U 12TB Storage

- SAN Ready NAS Storage Appliance (Server) with upto 24 x 500GB Hard Drive
- Dual AMD Opteron™ Processors 244+
- upto 4 x 2G Fibre Channel Ports for SAN
- upto 8 Gigabit Ethernet Ports for NAS
- Supports UNIX, Linux, FreeBSD, Windows
- 24 x Hot Swap Drive Bays for SATA or SCSI Drives
- 950W 3+1 Redundant Hot Swap Power Supply
- Custom Configuration is available
- Remote Support Available

6TB NAS starts at **$7,999**
12TB SANAS at **$14,999**

AMD64 Investment Protection
Migrate to 64-bit platforms seamlessly.
Add 64-bit application as necessary.

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

AMD 64 Opteron™

888.765.9686
www.Polywell.com/us/LJ

POLY SYSTEMS

AMD64 architecture reduces I/O bottlenecks, increases bandwidth, and reduces memory latency. Critical information gets to those who need it quickly and efficiently.

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

Opteron and ATHLON are trademarks of Advanced Micro Devices, Inc.. Quadro, nForce and Nvidia are trademarks of NVIDIA Corporation. All other brands, names are trademarks of their respective companies.

Make Stunning Schenker Graphs with GNU Lilypond

GNU Lilypond provides an easy-to-use, yet extremely powerful, tool for generating musical notation, including Schenkerian Analysis graphs.

BY KRIS SHAFFER

In the early twentieth century, Heinrich Schenker developed a method of analyzing tonal music that ties a piece's melody, harmony and form to a simple underlying musical idea. To illustrate his theory, he created a notational system that clearly depicts these relationships. Schenkerian Analysis, as it is called today, is a staple of music theory, but it is notoriously difficult to notate using the industry-standard, proprietary music notation applications Finale and Sibelius.

The Open Source world, however, has an excellent music typesetter in GNU Lilypond, which now runs natively on Linux, Mac OS X and Microsoft Windows. Lilypond not only produces beautiful sheet music, it also puts a great deal of control at the user's fingertips. Additionally, its text-to-music rendering method makes it easier for a typesetter to control hidden elements. This makes Lilypond a powerful tool for creating Schenkerian notation graphs, which—by their nature—require extreme control of positioning, as well as the masking and hiding of notational elements.

In this article, I cover the creation of a Schenkerian graph that contains all of the most common Schenkerian notational elements, with explanations of what each element signifies and the code required to produce it. I assume that the reader has at least a basic knowledge of Lilypond, and thus give instructions only for the nonstandard code used for Schenker graphs. I also assume that the user is using Lilypond 2.6, though most of the tools I cover are valid for any 2.x version of Lilypond. Armed with a working knowledge of Lilypond and with the techniques explained in this article, any user should be able to produce beautiful Schenker graphs—and some other forms of advanced musical notation—in less time, with less effort and difficulty than when using a graphical music notation application.

The Basics of Schenkerian Notation

There are a few simple steps to understanding a Schenker graph and how it represents an analysis of a piece. Two cardinal principles of tonal music form the foundation of Schenker's theory as an intrinsic part of the way we hear and perceive music. The first principle is the supremacy of the tonic (I) chord and the dominant (V) chord in the harmonic structure. That is, the chords built on the first and fifth notes of the scale. In the key of C major, this would be the C-major chord (I) and the G-major chord (V). The second principle is that the melodic structure is built upon a descending line, which ends on tonic (the first note of the scale).

A Schenkerian graph notates the structure of a piece in two main ways. First, rhythmic values are used to denote the structural importance of a note, not the length for which it should be played. Second, various musical markings—such as slurs, ties, beams and lines—are used to show the relationship of notes that have little structural importance to those that have greater structural significance. Schenkerian graphs also typically contain analytical markings such as Roman numerals for the harmony, scale-degree numbers and occasionally figured bass and analysis brackets.

As an example, let's use an excerpt of an analysis of J.S. Bach's Organ Chorale Prelude *Wenn wir in hoechsten Noten sein*, from Gene Biringer's book *Schenkerian Theory and Analysis: A Bridge from Traditional Harmony, Counterpoint, and Form to Advanced Studies in the Analysis of Tonal Music* (unpublished, Lawrence University Conservatory of Music). I chose this example because it clearly illustrates many of the standard Schenkerian notation elements, and I have made a few slight modifications to the graph to demonstrate the notation more completely. For the complete Lilypond file for this graph, see the on-line Resources.

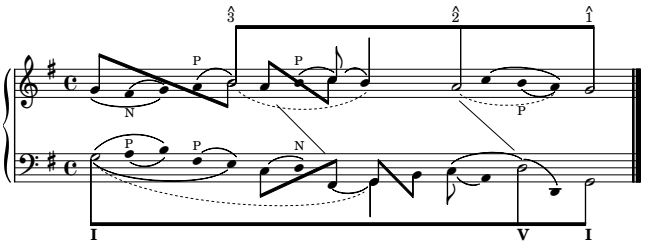

Figure 1. J.S. Bach: Organ Chorale Prelude *Wenn wir in hoechsten Noten sein*

In this example, note the use of different rhythmic values—half notes, quarter notes and eighth notes. In this case, as in most Schenker graphs, the half notes are the notes of the fundamental structure, and they are also beamed together to highlight the structure most clearly. Next, observe the use of ties, beams and slur marks in the graph. Slurs are used to connect notes of lesser structural significance with the fundamental structure. In the above example, the second and third notes in

DON'T BE SQUARE! GET CUBED!

L Series Laptop - LS1250-L
Light &Thin Performance
Starting at $1,562.54

G Series Laptop - GW1550-L
Essential Technology on a Budget
Starting at $1,302.54

X Series Laptop - XW1550-L
Extreme Technology & Performance
Starting at $1,608.04

R Cubed Technologies has provided pre-installed Linux laptops without OS tax since 2003. We customize the Fedora Linux distribution for each laptops configuration providing support for: PCMCIA, USB, FireWire, X, CD/DVD/CDRW/DVDRW, Sound, Power Management, Ethernet, Modem, Wireless, and more. Our laptops are equipped with Intel Centrino Mobile Technology. We also offer Windows dual boot options. All of our laptops come with a one year parts and labor warranty. Visit us online at www.shoprcubed.com or call 309.34.CUBED for details.

R3 Technologies™
Working Hard to Bring Technologies to Life.

309.34.CUBED
www.shoprcubed.com

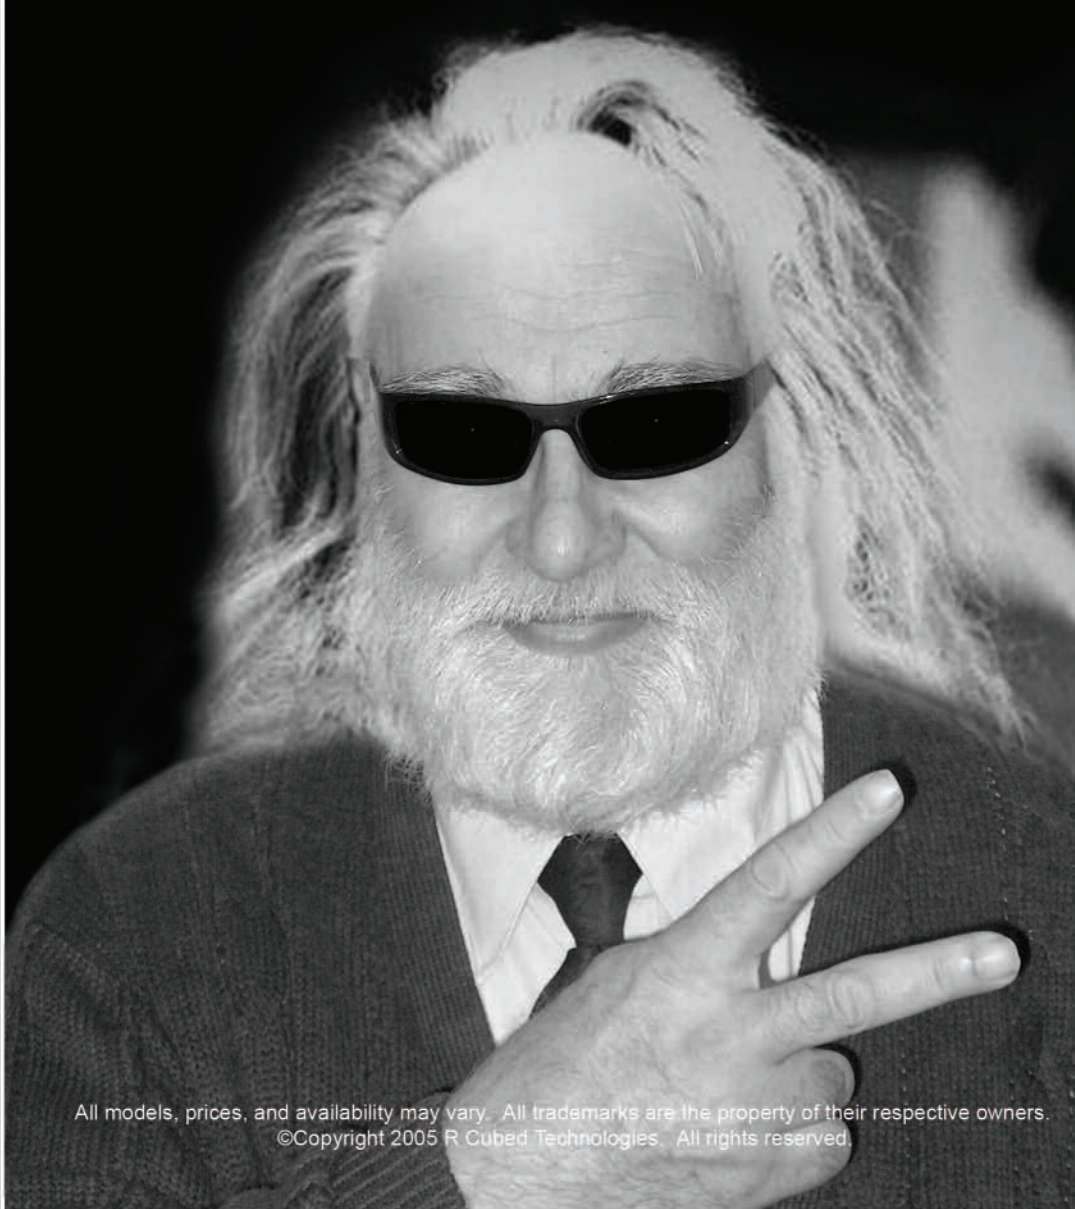

All models, prices, and availability may vary. All trademarks are the property of their respective owners.
©Copyright 2005 R Cubed Technologies. All rights reserved.

the upper staff are slurred showing that the F-sharp is a neighbor tone (marked by N) to the more structurally important G. The tie between the two Gs surrounding that F-sharp shows that the second G is a prolongation of the first G. Dotted slurs and ties are also used by some theorists to show extended prolongation of a note. In Figure 1, three dotted slurs or ties show extended prolongations of notes with other forms of embellishment between.

Lastly, observe the diagonal lines between the two staves. These lines are used to connect a melodic note and a bass note that coincide structurally, but they are not performed simultaneously in the piece. When examining the graph in Figure 1, one can see that every note in the example can be connected via slurs, ties or beams to the fundamental structure of the piece, thus showing the role of every note in the structure of the piece.

Creating a Lilypond Template for Schenkerian Graphs

Setting up a Lilypond file for a Schenker graph is fairly straightforward. A typical graph contains one or more grand staves, or piano staves, so one will likely begin with a piano template. To modify a piano template for a Schenker graph, add a few lines of code. Inside the PianoStaff brackets, but outside the individual Staff context brackets, add these lines:

```
\set Score.timing = ##f
\set PianoStaff.followVoice = ##t
```

The first line creates an unmetered score, with no barlines—typical for Schenker graphs. The second line is explained later.

Inside each Staff context, and inside \relative brackets if you use them, insert:

```
\override Staff.NoteCollision
  #'merge-differently-headed = ##t
```

This allows you to combine several layers of hidden voices—an important tool—without altering the note spacing.

The last global element is adding `raggedright = ##t` to the \layout section. I think this just looks better, but it also ensures consistency of measurement if you make significant edits to the graph after you've begun working on spacing. The piano template is now ready to be used for a Schenkerian graph. For an example template built upon a single grand staff, see the Resources.

Building a Schenkerian Graph in Lilypond

The first and most important part of the process of building a Schenkerian graph is to sketch the final graph by hand. The more complex the notation, the more valuable this will be. And even though Lilypond makes it easier than Finale or Sibelius to edit graphs after they have been created, you will still save much time and effort by sketching the complete graph by hand before typesetting it. It is also a good idea to mark off the beats that will be used. Because almost every voice in the Lilypond file will contain a number of skipped beats, it is essential to know the number and layout of beats ahead of time. One quarter-note beat for each notehead should suffice.

The Fundamental Structure

The next step is to typeset the fundamental structure, the half notes connected by an eighth-note beam. I chose to use two voices for each staff when creating this. One voice contains eighth notes with invisible noteheads, and the other contains half notes with invisible stems. The fundamental structure for the upper staff of the Bach graph looks like this:

```
<<
  {
    \override Beam  #'positions = #'(8 . 8)
    \override NoteHead #'transparent = ##t
    s1 b8[ s4. s1 a8 s4. s2 g8] s4.
    \revert Beam #'positions
    \revert NoteHead #'transparent
    }
\\
  {
    \override Stem #'transparent = ##t
    s1 b2 s1 a2 s2 g2
    \revert Stem #'transparent
    }
>>
```

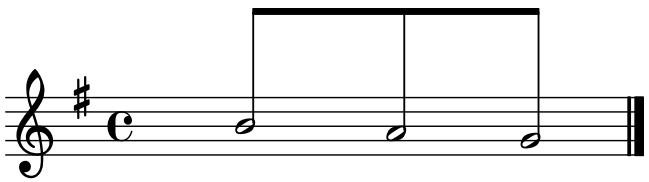

Figure 2. The fundamental structure uses half notes connected by an eighth-note beam.

Notice first that I override the beam positions, to make it level and out of the way of any notes and stems that may be placed under it. Also notice the use of the transparent property, one of your best friends when creating a Schenker graph in Lilypond. And of course, notice that the beats correspond exactly, with an eighth note and dotted-quarter skip in voice one, corresponding to each half note in voice two.

If you want to put scale-degree marks on each note, as in the Bach graph, Lilypond 2.6 now makes it possible without using LaTeX. You simply mark up the note like this:

```
b8[^\markup { \override #'(baseline-skip . 0.5)
\column { \small {^ 3} } }
```

The baseline-skip override should align the carat tightly over the numeral.

You also may notice that I chose to create multiple voices with brackets and back-slashes, rather than \voiceOne, \voiceTwo and so on. In my experience, the brackets are quicker, they make it easier to insert and delete voices, and they are less likely to cause alignment problems between voices.

Slurs and Layers

Once the fundamental structure has been created in each staff, next comes the surrounding notes. I typically begin with plain noteheads in one voice—separate from the two voices already created—and add beams, stems, slurs and

I CAN'T GET MY DATA OUT OF OUR CRM SYSTEM.

CUSTOMIZING OUR CRM SYSTEM IS OFF-LIMITS.

OUR CRM SYSTEM IS SO INFLEXIBLE.

SugarCRM™ gives you control over your CRM implementation that proprietary, closed source CRM restricts. Now you can fully manage your business critical data.

Start using the most powerful commercial open source CRM solution built on the modern LAMP architecture for managing sales, marketing, support and group collaboration activities today.

Get hooked on SugarCRM. Visit: **www.sugarcrm.com/linuxjournal** or call +1 408.454.6941.

Commercial Open Source Customer Relationship Management

SUGARCRM™

www.sugarcrm.com
www.sugarforge.org

1.87SUGARCRM
+1.408.454.6941

Copyright © 2005 SugarCRM, Inc. All rights reserved. SugarCRM and the SugarCRM logo are trademarks of SugarCRM, Inc. in the United States, the European Union and other countries.

additional voices as I continue. This helps me better organize my code. To create a voice with plain noteheads, begin the voice with the code:

```
\override Stem #'transparent = ##t
\override Stem #'length = #0
```

and follow with all quarter notes or quarter-note skips, never rests. The second line becomes useful when adding slurs. Because slurs attach to the stem if the stem and slur are on the same side of the note, you can use this line to cause all slurs to attach only to noteheads. Just remember to cancel it out when you add a visible stem. When you reach the end of the voice, remember to \revert anything you \override.

Once the notes are added, you can begin adding the necessary slurs, beams and other appropriate marks. Multiple layers of slurs are intrinsic to Schenkerian notation, but they can be cumbersome in Lilypond code. There are two ways to accomplish it. The first is to use the phrasing slur tool. This allows you to create a lower layer of slurs with (and) and an upper layer with \(and \). This allows for only two layers of slurs, but it does let you keep both layers in the same voice. If two layers of slurs are all you need, this may help you keep your code cleaner and save you a little work.

If you need more than two layers—note the four layers of slurs on the first note in the Bach example—you must create multiple voices. If you require three layers of slurs, create three voices. In the first voice, begin with:

```
\override Stem #'transparent = ##t
\override Stem #'length = #0
```

as before, and follow this with all the notes in the line (and the skips and \revert commands). Insert the first (lower) layer of slurs in this voice.

In the second voice, begin with:

```
\override NoteHead #'transparent = ##t
\override Stem #'transparent = ##t
\override Stem #'length = #0
```

and follow with all the notes and the second layer of slurs. This attaches each slur to an invisible note in the same place as the visible notehead from voice one. If you want, you can replace the unneeded notes in this voice with skips, but it is unnecessary. The third voice will look like the second voice, but it will include only the third layer of slurs.

After making a few minor spacing adjustments, your code may look something like this (a variation of the beginning of the lower staff of the Bach example):

```
<<
  {
    \override Stem #'transparent = ##t
    \override Stem #'length = #0
    \once \override TextScript #'extra-offset =
      #'(-11 . -2.5)
    g4 a( b) fis( e)
    \revert Stem #'transparent
    \revert Stem #'length
    }
\\
  {
    \override NoteHead #'transparent = ##t
    \override Stem #'transparent = ##t
    \override Stem #'length = #0
    \once \override Slur #'extra-offset =
      #'(0.5 . 0.75)
    \once \override Slur #'height-limit = #1.5
    g4( a b) fis e
    \revert NoteHead #'transparent
    \revert Stem #'transparent
    \revert Stem #'length
    }
\\
  {
    \override NoteHead #'transparent = ##t
    \override Stem #'transparent = ##t
    \override Stem #'length = #0
    \slurDown
    \once \override Slur #'extra-offset =
      #'(-1.25 . 0)
    \once \override Slur #'height-limit = #2.75
    g4( a b fis e)
    \revert NoteHead #'transparent
    \revert Stem #'transparent
    \revert Stem #'length
    }
>>
```

Figure 3. Using layers of slurs helps you organize your code clearly.

Even when using only two layers of slurs, I prefer this method rather than using slurs and phrasing slurs combined. It gives me the same method in every graph, it organizes my code more clearly and when I edit slur properties, I always use the same commands for any layer. Otherwise, I would alternate between overriding slur properties and phrasing slur properties.

Editing and Tweaking Slurs

When using slurs in Schenker graphs—especially when using multiple layers—you likely will need to edit some of the slur properties in your graph. The simplest edits are \slurUp and \slurDown, which cause the following slur to be created above or below the notes, respectively, and \slurDashed, a new tool in Lilypond 2.6, which creates a dashed slur.

Another common tweak I find useful is:

```
\once \override Staff.Slur #'height-limit = #x
```

This allows me to specify how deep or shallow the slur should be drawn (represented by the value x), and it is especially helpful for layered slurs or for slurs under and over text. Occasionally, I have to specify the entire set of coordinates for a slur manually. This lets you create some funky slurs, which

some Schenkerian theorists love to use. For that, use the following code, where x and y are horizontal and vertical distances, measured in staff lines:

```
\once \override Slur #'control-points =
  #'((x . y) (x . y) (x . y) (x . y))
```

And, as with just about any notational element in Lilypond, you can alter a slur's extra-offset property, moving the entire slur without altering the shape:

```
\once \override Slur #'extra-offset = #'(x . y)
```

See the example of slur layers in Figure 3 to observe extra-offset and height-limit in action.

Cross-Staff Diagonal Lines

Occasionally, a melodic note corresponds to a bass note harmonically, but they are not sounded simultaneously and thus are not aligned vertically in the score. In Schenkerian notation, a simple diagonal line connecting the notes suffices to make this connection. Unfortunately, such a line is not as easy to create in Lilypond as in a graphical editor. However, it can be done rather painlessly with `\change Staff`. When creating our template, we added the line:

```
\set PianoStaff.followVoice = ##t
```

to our file. That line combined with `\change Staff=LH` or `\change Staff=RH` creates a diagonal line that follows the voice from one staff to the other. Thus, if you create a new voice in the upper staff with the following code:

```
\override Stem #'transparent = ##t
\override NoteHead #'transparent = ##t
\override Stem #'length = #0
s1 s4 e4 s
\change Staff=LH
fis,4 s2
\revert Stem #'transparent
\revert NoteHead #'transparent
\revert Stem #'length
```

you will get the first diagonal line in the Bach example, descending from the upper staff to the lower staff. The transparent noteheads and stems cause Lilypond to render only the diagonal line. Using invisible notes also allows you to alter the pitch of the start and end notes to adjust the height of each end of the line. Though this may seem to be overkill, the entire block of code easily can be cut and pasted to another voice or file, with the necessary adjustments being only height and beat placement, making this an easy solution. (If you really want to click and drag the line onto the graph, open the finished graph in an image editor and add the line there.)

The Unfolding Symbol

The last Schenkerian idiom I cover here is the unfolding symbol. Briefly, this symbol signifies a harmonic connection between two notes in a melody. They typically occur in pairs, showing the use of two concurrent harmonic voices

ASA COMPUTERS

Want your business to be more productive?

The ASA Servers powered by the **Intel® Xeon™ Processor** provides 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.)

6TB + in 5U—$8,450
Intel 7501, Dual Intel® Xeon™ 2.4GHz
512 MB DDR ECC RAM *Max: 8GB*
6TB + IDE Storage
Dual Gigabit LAN, CD+FD, VGA
Options: SATA Drives, Firewire, DVD+RW, CD+RW, 64 Bit OS Configurations, etc.

1U Dual Itanium IDE—$3,925
Dual Intel® Itanium®2 1.4 Ghz
2 GB ECC DDR
1 of 4 x 40 GB HDD
Dual Gigabit LAN
Based on Supermicro 6113M-i

14" Deep Appliance Server—$865
Intel® Xeon™ 2.4 Ghz Processor
40 GB Hard Drive, One GigE
Options: CD, FD, 2nd HD, Your Logo on Bezel
Call for Low Cost Options.

1U Dual Xeon™ EM64T Superserver—$1,799
SuperMicro 6014H-82 Barebones
1 of 2 Intel® Xeon™ 2.8 GHz 800 FSB
1 GB DDR II-400 RAM *Max: 16GB*
36 GB 10K RPM SCSI *Max: 4 HS HDD*
CD+FD, Dual GigE, VGA, RAILS
Options: RAID, etc.

Your Custom Appliance Solution
Let us know your needs, we will get you a solution

ASA Colocation
$50 per month for 1U Rack - 20 GB/month

ASA Colocation Special
First month of colocation free.*

Storage Solutions
IDE, SCSI, Fiber RAID solutions
TB storage options
3Ware, Promise, Adaptec,
JMR, Kingston/Storcase solutions

Clusters
Rackmount and Desktop nodes
HP, Intel, 3Com, Cisco switches
KVM or Cyclades Terminal Server
APC or Generic racks

All systems installed and tested with user's choice of Linux distribution (free). ASA Colocation—$50 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 inside XEON™

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 typographical errors.

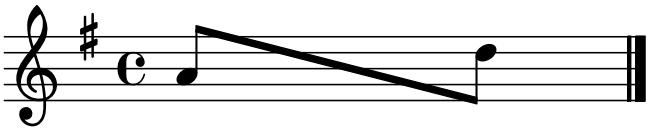

Figure 4. The unfolding symbol shows a harmonic connection between two notes in a melody.

in one melodic line. They are surprisingly easy to create. When two simultaneous notes in a line are to be connected with the unfolding symbol (as in the lower staff of the Bach example), one simply needs two notes connected by beaming brackets, with the commands \stemUp and \stemDown in the appropriate locations. Of course, one must remember to remove stem transparency before creating the unfolding symbol and insert eighth-note skips appropriately to preserve vertical alignment:

```
\override Beam #'positions = #'(1 . -4)
\stemUp
g8[ s
\stemDown
b8] s
```

Notice the use of beam positions to adjust the height of the stems and the beam angle. When other notes occur between the two notes to be connected with an unfolding symbol, as in the upper staff of the Bach example, put the unfolding notes in one voice and the independent noteheads in another, with appropriate skips in each voice. For example, if the first voice contains:

```
\override Beam #'positions = #'(3 . -2.5)
\stemUp
a8[ s s2
\stemDown
d8] s
\revert Beam #'positions
```

and the other contains:

```
\override Stem #'transparent = ##t
s4 b c s
\revert Stem #'transparent
```

the end result will turn out like Figure 5.

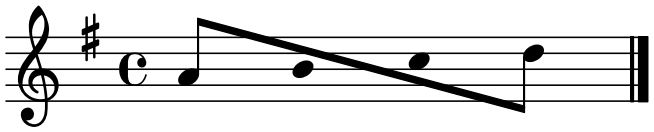

Figure 5. Other notes can appear between two notes connected with an unfolding symbol.

Conclusion

Creating Schenkerian graphs in a graphical editor like Finale or Sibelius is enough to make many theorists revert to pencil and paper. The process is long and difficult, making changes to finished graphs is nearly impossible and you must do the same things to each graph every time you create a new one. However, with GNU Lilypond and the above tools, any musician can create beautiful Schenker graphs with minimal headaches and maximum control. Lilypond's text-to-music method makes it easy to edit hidden elements, modify finished graphs, and cut and paste code to future projects. Though the methods take time to learn, in the long run Lilypond saves time, energy and frustration, all the while creating stunning output. The tools and examples in this article should put you well on your way to creating beautiful Schenker graphs and some other forms of advanced musical notation with this great application.

Resources for this article: www.linuxjournal.com/article/8583.

Kris Shaffer lives in New Haven, Connecticut, where he is pursuing a PhD in Music Theory at Yale University. An open-source enthusiast as well, he has written for Linux.com, NewsForge.com and OSNews.com. Kris is also co-founder of **www.AmSteg.org**, an on-line community for composers and music theorists, which is making its debut in Fall 2005. His personal Web site is **www.shaffermusic.com**.

ospi
Open Source Professional Institute
Preparing Business for Tomorrow

Linux Certification

Train Now for LPIC-1 Linux Certification
Not for the faint of heart...
Obtaining Linux Certification will set you apart from the multitude of Linux IT professionals and prove to potential clients and future employers that you've got what it takes.

~ 2006 ~
1/16 Pittsburg
1/30 Memphis
2/13 Miami
2/27 NYC
3/13 Detroit
3/27 Las Vegas

LPIC-1 Training
Mon – Fri 8am – 4pm

Additional Offerings:
Building an Enterprise Mail Server
2 days 5 – 7pm
Building an Enterprise Web Server (LAMP)
2 days 5 – 7pm

Call NOW to Reserve Your Space
1-800-316-7912
www.ospinstitute.com

Scale up performance. Scale down costs.

Appro lets you scale up performance while scaling down cost. The choice and the cost savings are yours.

Appro *HyperBlade* Clusters

Full-Cluster
Up to **80** nodes

Mid-Cluster
Up to **50** nodes

Mini-Cluster
Up to **17** nodes

- Flexible, modular and scalable architecture
- Blade servers fit all cabinet sizes
- Dual AMD Opteron™ processors per node
- Up to 16GB ECC DDR Memory per node
- One PCI-X slot per node
- Up to two internal HDDs
- Remote Management Options
- High-Speed Interconnect Options: Myrinet, Dolphin, Quadrics and InfiniBand™

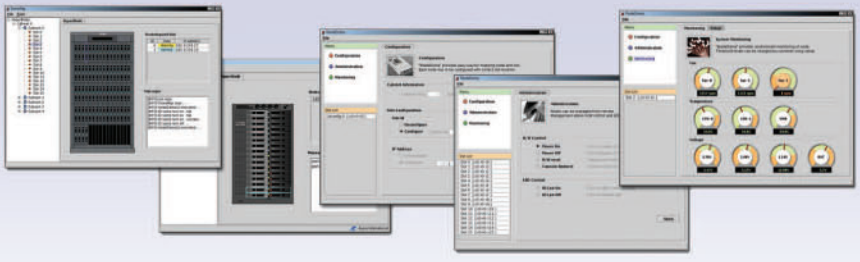

✔ *Specially created for the Appro* Hyper*Blade servers*
✔ *Outstanding hardware and software management tool*

Appro **BladeDome** Remote Server Management

- GUI & command line interfaces
- In-band and out-of-band control
- Remote reset and power cycle
- Platform Monitoring: fan fail, over-temperature & voltage
- Failure alert e-mail notifications
- Enhaced security
- Multiple user account set-up
- BCC redundancy

AMD Opteron™ Processors
- Integrated AMD HyperTransport™ technology allows for concurrent multiple processors in a single system.
- Shorten run-time cycles and increase bandwidth for processing computing requests.
- 32 bit applications while you migrate to 64 bit computing for long-term investment protection.

APPRO
HPC Cluster Solutions

Appro has everything you need to create a network blade cluster-ready. On-site maintenance and installations services are also available.

For more information, please visit www.appro.com or call Appro Sales at 800.927.5464, 408.941.8100.

AMD 64 Opteron

DVD Mastering Using QDVDAuthor

Dan Sawyer, an expert in Linux-based video production tools, shares the inside scoop on mastering DVDs using QDVDAuthor.

BY DAN SAWYER

When I started using Linux seven years ago, I was told repeatedly that only a rank moron would make a serious attempt to do video and multimedia production on a raft load of open-source UNIX tools. At the time, I was just starting into the world of digital production and had a comfortable set of Windows workstations equipped with appropriate tools, and I had no compelling reason to move to Linux, save one: I was sick of Windows.

Consequently, in the intervening years, I have been gradually moving my studio from Windows to Linux. Linux has long been capable of editing digital video with applications like Kino and Cuisine, and of capturing and editing other video with Cinelerra and its predecessors. Even longer, Linux has been an excellent and capable 3-D graphics creation and rendering platform, and The GIMP has long been adequate for most raster graphics manipulation needs. Two areas in which Linux has lagged behind the competing operating systems are in compositing and complex DVD authoring. Although the former continues to be a problem, the latter is beginning to come into its own.

Recently, I had the opportunity and the spare time to revisit the field of Linux DVD authoring and was pleasantly surprised by what I found. Although the command-line tools for DVD authoring have long had the technical capability to construct a complex menu-driven DVD, the GUI tools suitable for use by end users and artists have been sorely lacking. I was pleased to find that this situation has changed. The task: to put together a promotional reel for my independent film, packaged on an attractive, menu-driven DVD.

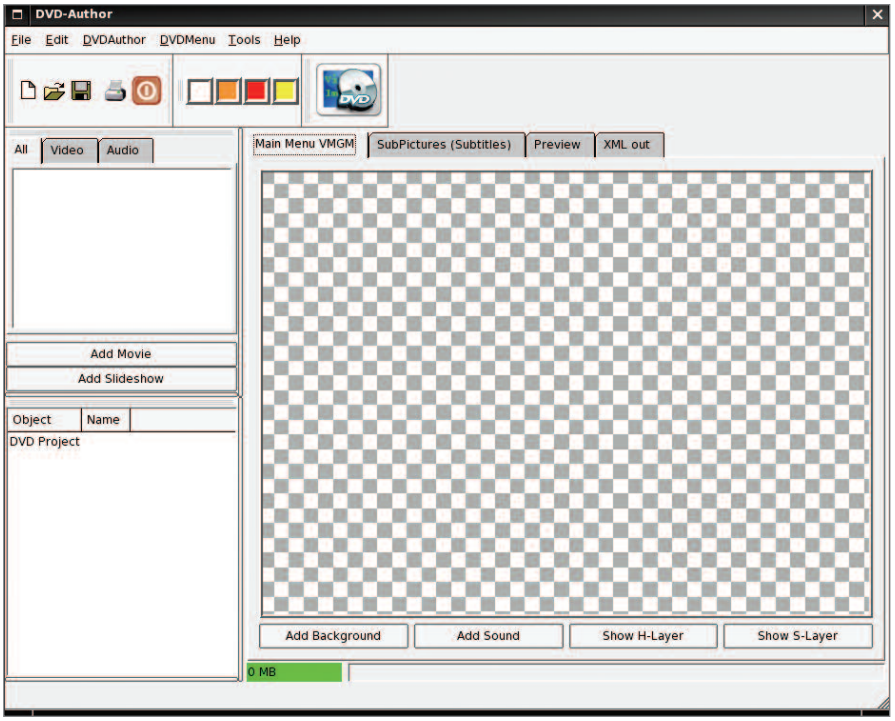

Figure 1. QDVD Main Window

After browsing around through the various available packages, I chose QDVDAuthor. On balance, it offered the best compromise of smooth workflow, content control and minimal dependency headaches. Its current major weakness (namely, the obscure manner in which it gives direct access to the DVD file structure for hacker-level tweaking) was not of concern to me as I was not looking to put together a commentary track or hide easter eggs.

As we embark, I hardly need to remind you, this is beta-level software. *Always* save your work; QDVDAuthor is crash-prone in some of its less-well-developed areas, and you don't want to be caught unawares.

After installation (see the sidebar for program and dependencies info), we open up to the main window (Figure 1).

To start a new project (File→New), we need to run through the setup wizard (which appears when you specify the creation of a new project), specifying temp directories, project title and data destination directory (Figures 2 and 3).

Only two screens long, the wizard is quickly finished and we're ready to start building the DVD itself. For the purpos-

Why is LPI the Global Standard in Linux Certification?

Trusted.

All Linux Professional Institute certification programs are created using extensive community input, combined with rigorous psychometric scrutiny and professional delivery. We test the whole continuum of important Linux skills - we don't just focus on small, subjective tasks. LPI exams are not simply an afterthought used to help sell something else. LPI is a non-profit group that does not sell software, training or books. Our programs and policies are designed to meet educational requirements, not marketing.

Accessible.

LPI exams are available in seven languages, at more than 7,000 locations, in more than 100 countries. You take LPI exams when you want, where you want. In addition, special exam lab events around the world make our program even more affordable. And because we don't make exclusive partnerships, LPI is supported by a broad range of testing centers, book publishers and innovative suppliers of preparation materials.

Independent.

You switched to Linux to get away from single-vendor dependence. So why trade one form of vendor lock-in for another? LPI's program follows the LSB specification, so people who pass our tests can work on all major distributions. Because of its strong grass-roots base and corporate support both inside and outside the world of open source, LPI goes beyond "vendor-neutral" to truly address community needs.

LPI is IT certification done RIGHT!

For more information, please contact us at Info@lpi.org or visit us at www.lpi.org.

Linux Professional Institute

Figure 2. New Project Wizard—Defining the Temp Drive

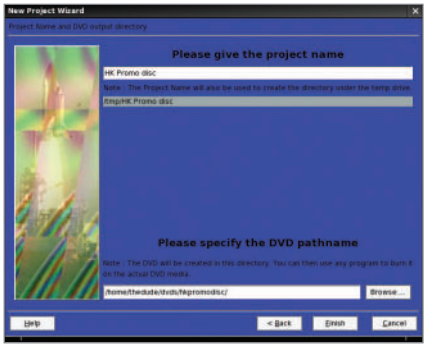

Figure 3. New Project Wizard—Specifying the Project Name and DVD Pathname

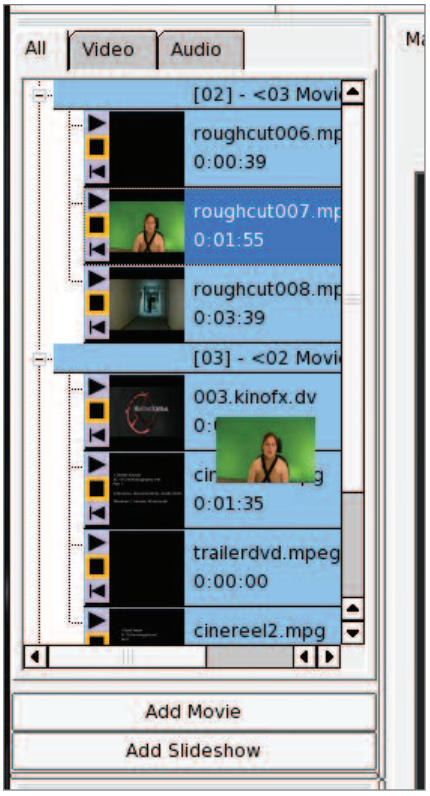

Figure 4. Dragging and Dropping Files

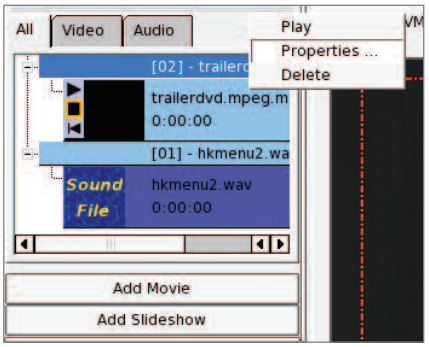

Figure 5. Reordering Tracks

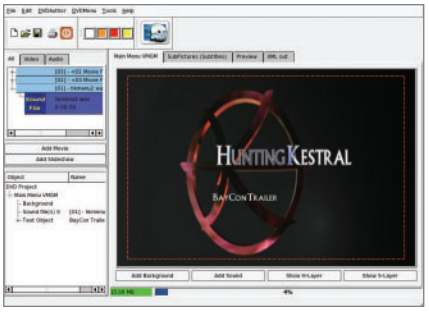

Figure 6. Menu Background and Soundtrack Loop

es of this tutorial, I'm assuming you already have DVD-spec MPEG-2 video to work with. If you don't, see the sidebar for a quick run-through on encoding DVD video. I'm going to be building a DVD with the trailer and blooper reel from my independent film, *Hunting Kestral*, which was edited in Kino with titling done in Blender.

Now that our project is set up, let's begin by importing the videos we want to use. Assuming that you've proactively broken up your video to avoid the sync problem, you're going to want to import your video files so that they're organized the way you want them to be and (most important) playable simply from one button with a minimum of fuss. To do this, click the Add Video button to the left of the workspace in the main window, and then select your videos. Selecting multiple files in the import dialog (Ctrl-click) imports the files into a single title heading. If you later want to add more files to a particular title, simply import them and then drag and drop them between titles in the All tab in file list window (Figure 4). By dragging and dropping, you can also mix up the order in which your titles play. For more judicious control over the order in which the files in a title play, right-click on the title and select Properties (Figure 5), and reorder the tracks using the up and down buttons.

Once you've finished importing all your footage, it's time to begin constructing your menus. You'll want to build your menu structure from the end up, starting with the deepest menus and working your way back up to the main menu, so that when you're linking up all the buttons, you don't have to retrace your steps. I'm going to proceed with the main menu. Because you'll use the same set of steps and tools for creating every menu, feel free to expand as far as you like. In the main window, click the Add Background button. From here, you can select any compliant still image or video (MPEG-2) file that you'd like to underlay your main menu.

If you want to have a looping soundtrack running in your menu, click the Add Sound button, and load in any compliant audio file (.wav, .mpa, .mp3, .mpega, .ogg and so on). However, there is a bug here to beware of: QDVDAuthor seems to have a curious oversight in its design, in that it imports most anything for sound, but it won't actually encode the audio to a suitable format for multiplexing. The transcoding dialog (which you get to by right-clicking on a clip and selecting Properties) seems to crash for no good reason when working on audio, and as such you can't automatically transcode the audio when you're outputting the final project. This means that, at the moment, audio loops for your menus have to be in an mplex-friendly format (AC3 or MPEG-Audio). To get suitable audio out of your .wav, .ogg or .mp3 file, simply type the following at your command prompt:

```
ffmpeg -acodec mp2 foo.mp2 -i foo.wav
```

Once that's done, you can import the .mp2 file into QDVDAuthor and proceed normally.

Figure 6 shows the menu background, which I made in Blender, and the soundtrack loop loaded up.

Now that the background is set up, it's time to construct the buttons. Because this is an introductory article, we work only with text buttons, although QDVDAuthor is perfectly capable of video buttons or image buttons as well.

O'REILLY®

EMERGING TELEPHONY CONFERENCE

JANUARY 24–26, 2006
SAN FRANCISCO AIRPORT MARRIOTT
conferences@oreilly.com/etel

THE FUTURE IS CALLING

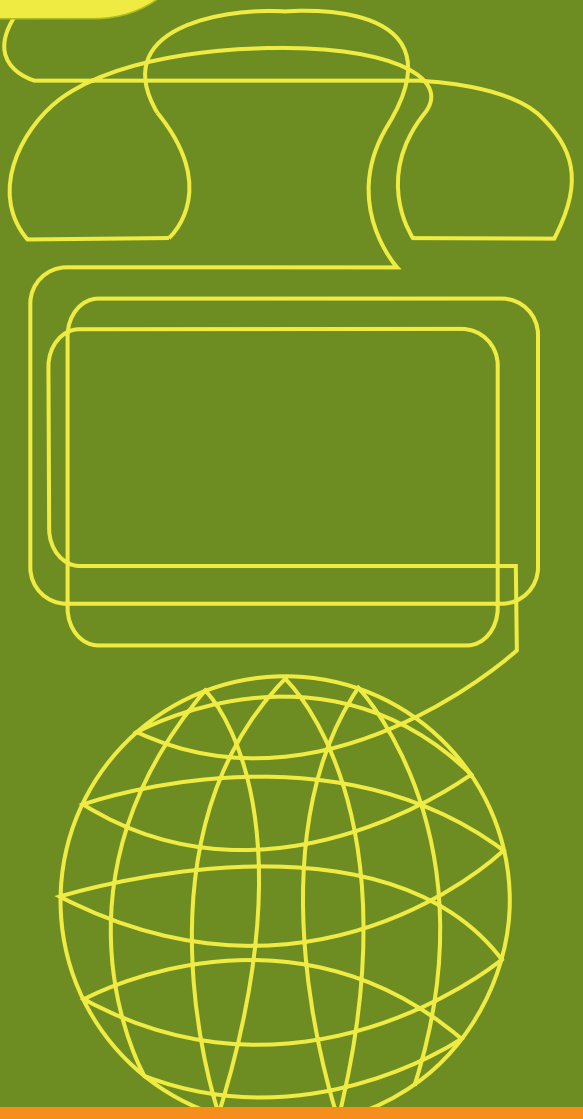

ETel prepares you to ride the wave of disruption and improve your communications in unconventional ways.

From projects to products to players, ETel provides a map of the evolving telephony horizon and gives you the charts you need to navigate the new communications opportunities ahead.

REGISTER BY DECEMBER 5 AND SAVE $300.00
conferences@oreilly.com/etel

© 2005 O'Reilly Media, Inc. The O'Reilly logo is a registered trademark of O'Reilly Media, Inc. All other trademarks are property of their respective owners. #50388

VIDEO ENCODING

If you're working from footage taken on a miniDV camera, outputting a DVD-ready video from Kino is relatively simple. If you haven't used it before, Kino's interface is comfortable and easy to navigate. Good user guides are available at several places on the Web, most obviously at Kino's home page (kinodv.org/article/archive/13). It's not a multitrack editor, more's the pity, but for quick-and-dirty edit work with basic transitions and soundtrack mixing, it works encouragingly well. When outputting video from Kino, I've found that I get the best results (for both video quality and a minimum of sound sync slippage) with the dual-pass encoding in the DV Pipe screen.

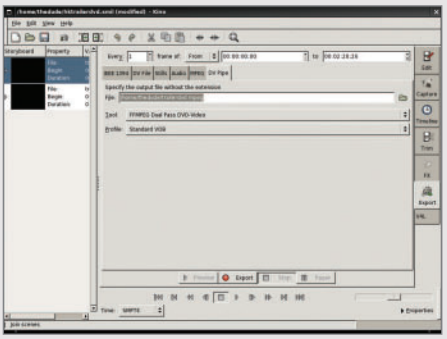

DV Pipe Screen

On the off chance that you're wanting to burn DVDs from your PVR, you still need to get the files into the right format. Mencoder is great for this, though it has a confusing array of options. Here's a sample command argument for moving from xvid to DVD-compatible MPEG-2:

```
mencoder -ovc lavc -lavcopts vcodec=mpeg2video -oac lavc -lavcopts \
acodec=mp2:abitrate=512 foo.avi -o foo.mpg
```

An important caveat about encoding to DVD-format MPEGs: every Linux video encoder I have ever run into uses FFmpeg or MJPEGTools as a back end, and they both have the same problem—a big one. They both seem to have a bug that causes a slip of sound sync progressively throughout the file, becoming noticeable after about the first two minutes of footage. It's a problem in the library that I've not found a way around, though it is markedly less pronounced using FFmpeg than MJPEGTools. This is the biggest and most troublesome hurdle still facing Linux DVD authors. The only solution I've found to this deeply irritating problem is to slice your video into two– to five-minute tracks and use each of these tracks as separate titles on your DVD. It's an ugly solution, and not the kind of thing you want to talk about at parties, but for the moment it's the best we can do. In an ideal world, the good folks who maintain these projects would fix the issue, but as this is a common problem for many commercial MPEG encoders, I'm not holding my breath. (I should add, dear reader, on the off chance that this is a user-brain-dead error and I'm missing something obvious, I look forward to your hate mail with cheerful enthusiasm.)

In case you want to strike out on your own with the available command-line tools (mencoder, FFmpeg and mjpegtools), here are the vital stats you'll need to encode a serviceable DVD video file (all numbers are for NTSC):

Video:

- 720x480 with 4:3 (standard) or 16:9 (anamorphic) aspect ratio.
- MPEG-2 @ up to 98,00kbps

Audio:

- 48khz @ 32–1,536kbps
- PCM, AC3, MPEG-1 Layer2
- Up to eight audio tracks encoded

To create a text button, right-click on the work space, and select Add Text. Doing so turns your cursor into a cross, with which you click and drag to draw the text box for your text to fill. Don't worry if you draw the wrong shape or put it in the wrong location, such mistakes are easily remedied after your text has been specified.

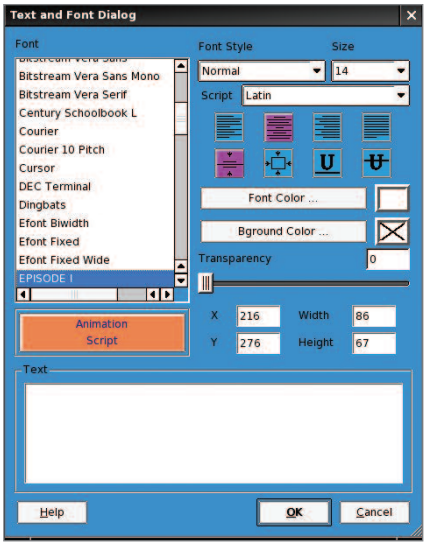

Figure 7. Text Creation Dialog

After drawing the box, the text creation dialog appears, and you can select from any of the fonts installed in your X11/fonts directory, as well as set the color, alignment, size, style and background color for your text box. Specialized TrueType fonts should be dropped into this directory before starting QDVDAuthor, and they will appear in the font selection dialog. This is also the place where you adjust the dimensions of your box and its placement (although placement can also be adjusted by simple drag and drop in the main window).

Once you've finished defining your

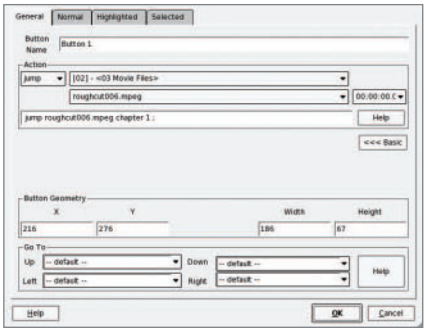

Figure 8. Defining the Text Button

text box, click OK, and in the main window, right-click on your text box and select Define as button (Figure 8).

Here, you assign the action that the button will take (jump to a file, call a menu or resume). If you click the Advanced>>> button, you can further define the way the navigation controls move the cursor around the DVD menu (the up, down, left and right list boxes), routing each directional button to a different button on the screen. So, for example, in my current project, I have five buttons, and I want the viewer to be able to navigate between each button in a fairly obvious fashion (Figures 9 and 10).

Figure 9. Current Project Showing the Five Text Buttons

Figure 10. Routing the Directional Buttons

As such, I set up the focal Play All button for simple direct navigation, as shown in Figure 11.

Figure 11. Setting Up the Play All Button

Figure 12. DVD Export Button

DVD FILE STRUCTURE

The difference between a data DVD and a video DVD is essentially the file structure and video format. The proper encoding of the file structure is handled by DVDAuthor, the back end on which all Linux DVD programs depend. It takes an XML file and builds the DVD image from it. Here is the DVDAuthor output from the project I built for this article:

```
<dvdauthor dest="/home/user/dvddirectory/" jumppad="yes" >
  <vmgm>
   <menus>
    <video format="ntsc" resolution="720x480" />
    <pgc entry="title" >
     <vob file="/tmp/HK Promo disc/Main Menu VMGM_menu.mpg"
pause="inf" />
     <button name="1" >jump title2; </button>
      <post> jump vmgm menu 1; </post>
    </pgc>
   </menus>
  </vmgm>
  <titleset>
   <menus>
    <pgc>
     <post> jump vmgm menu 1; </post>
    </pgc>
   </menus>
   <titles>
    <pgc>
     <vob file="/home/user/dvdmenu1.mpeg.vob" />
     <vob file="/home/user/trailerdvd.mpeg.vob" />
     <vob file="/home/user/video/cinereel1.mpg" />
     <vob file="/home/user/video/cinereel2.mpg" />
     <post> call vmgm menu 1; </post>
    </pgc>
   </titles>
  </titleset>
</dvdauthor>
```

Now, all that's left is getting the Play All button to work properly, which is deceptively easy. Since we grouped our videos all under one title, ordering them upon import in the order in which we wanted them to play, all that's needed to Play All is to link the button to the first chapter of the title. Once started, it will play through all videos in that title before returning to the main menu. More complicated playlist arrangements are possible and fairly straightforward, but are beyond the scope of this article.

Once your buttons are configured and your videos ordered, you're ready to burn. Click on the DVD Export button (Figure 12), and it'll bring up the export batch manager window (Figure 13).

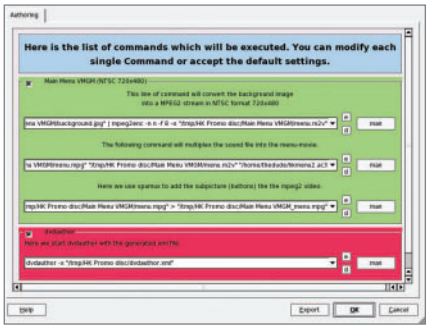

Figure 13. Export Batch Manager Window

Here, you will be able to make any final tweaks to your project before building the DVD—and burning, if you prefer to do it from here rather than from K3b.

Once you click OK, all of these

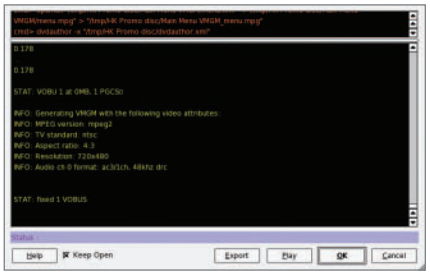

Figure 14. Executing the Commands

QDVDAUTHOR AND DEPENDENCIES

QDVDAuthor is available at **qdvdauthor.sf.net**.

It depends on the following packages: Transcode, PCM2AIFF, toolame, dvdrecord, dvdauthor, FFmpeg, mjpegtools, arecord, oggdec, lame, mplayer/mencoder, dvd-slideshow, sox, imagemagik, mkisofs, growisofs and dvd+rw-format.

commands are executed in order in the window shown in Figure 14.

Watch carefully for any error flags printed in red. If you ignore them, you'll find yourself burning coasters, if you have a burnable image at all. If you do get any red flags, re-run the batch, clicking the Export rather than the OK button, which will export a shell script that you can disassemble and tweak to track down and correct your problem. This isn't often necessary, but every once in a while the program just doesn't generate the proper output and it has to be hand-tweaked.

If all has gone well, you now have a simple, menu-driven DVD suitable for exhibit on any consumer DVD player and TV.

The authoring toolset available for Linux, although still having its rough patches, is now finally capable of producing professional and complex DVDs with audio commentaries, video menus, animations, photo galleries, score-only tracks, chapter-selection menus, playlists and easter eggs. Each of these specialized structures requires a bit of elbow grease, but all do work together. With a little poking around and the occasional XML tweak, Linux is finally up to the task of filling the authoring spot in the multimedia studio pipeline. Future improvements in QDVDAuthor and its competing programs are sure to make the situation even better.

Now, if only we could get an open-source compositor that's up to snuff....

Dan Sawyer is a freelance director/producer running the backbone of his small studio on Linux. He has been an enthusiastic advocate for Free and Open Source software since the late 1990s, when he founded the Blenderwars filmmaking community (**www.blenderwars.com**). Current projects include the independent SF feature *Hunting Kestral* (**www.blenderwars.com/kestralmannix**) and *The Psyche Project*, a fine art photography book centering on strong women in myth.

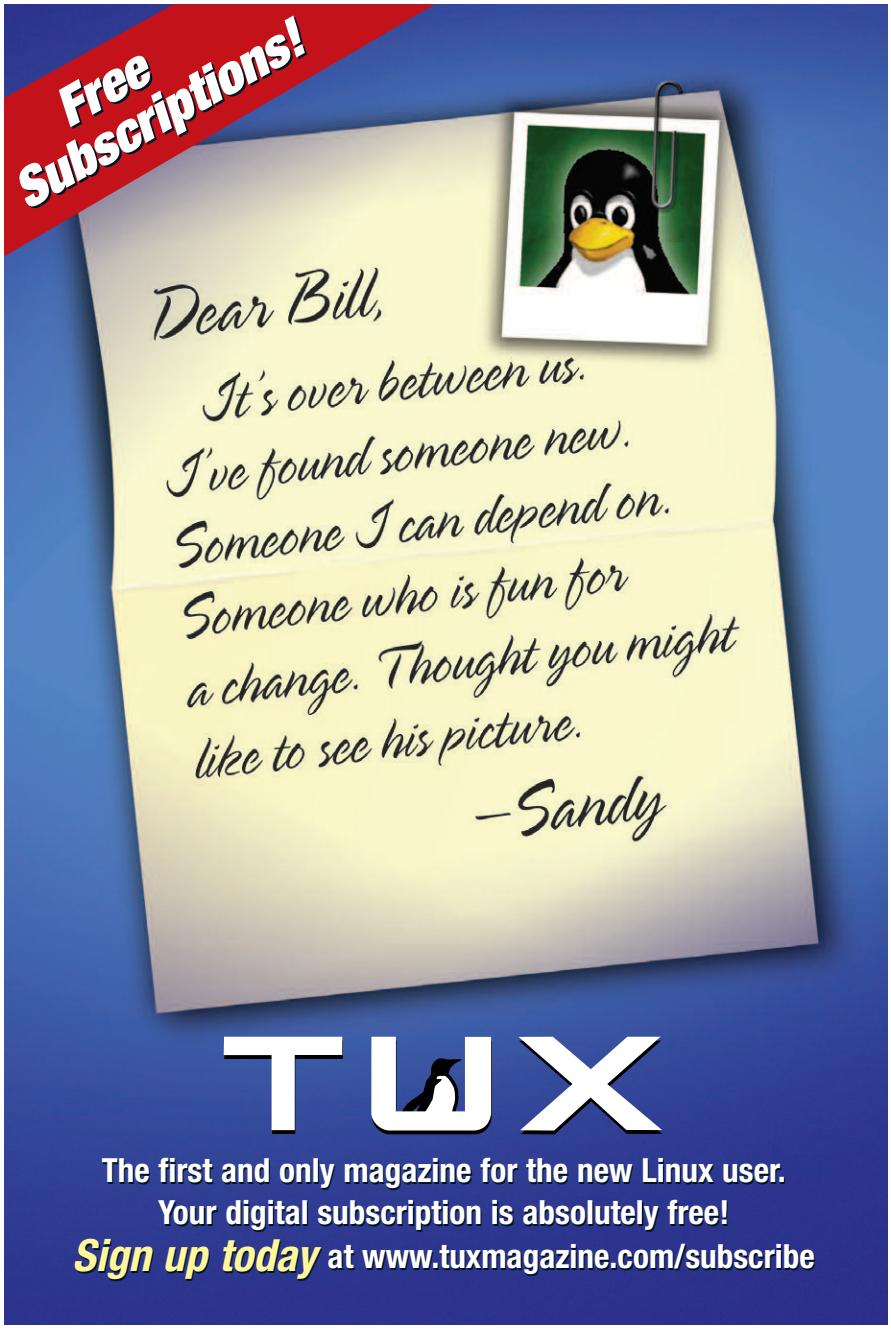

FREE SOFTWARE
FOUNDATION

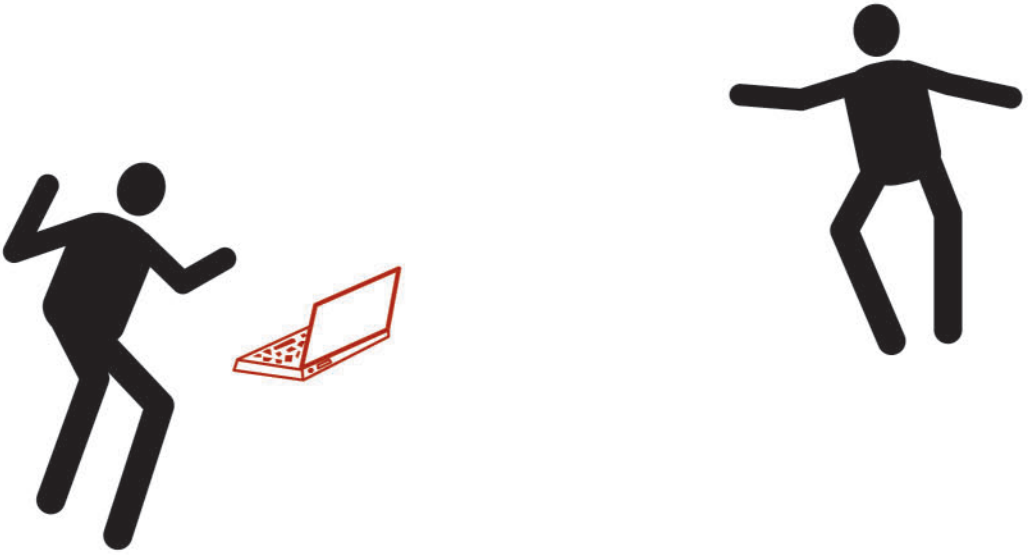

be FREE

Join us at member.fsf.org

A Linux DVR Is No Myth—It's MythTV!

James Turner gives us an overview of MythTV—a Linux-based TiVo replacement. BY JAMES TURNER

The advent of personal digital video recorders (DVRs) has transformed the experience of watching TV for millions of people. The VCR may have freed viewers from having to watch programs when the networks wanted them to, but the DVR has given them dramatically more freedom and control.

Most consumers use a set-top DVR, either buying it outright or getting it as part of a bundle with a satellite or cable package. But by their nature, these boxes fail to realize the benefits of a DVR fully. For one thing, because of the politics of business, DVR manufacturers have been reluctant to develop technologies to allow viewers to skip commercials automatically. Also, commercial systems require monthly subscription fees to receive the viewing guides, which can easily exceed the original cost of the recording in a single year. Finally, the manufacturers discourage owners from making simple modifications, such as adding additional disk space or networking their DVRs to allow other TVs in the house to watch recorded content.

This brings us to solution number two—building a DVR from scratch. Several packages are available for Windows—none for free—that do a competent job of providing DVR capabilities. But if you want a high-quality DVR that runs under Linux, MythTV is the way to go. This article walks you through the steps you need to set up MythTV on an already functional Linux system.

The architecture of a MythTV box is fairly simple. A dæmon process called mythbackend is responsible for actually talking to the tuner cards, figuring out what programs should be recorded and otherwise handling the day-to-day business of being a DVR. In theory, that's all you need to have running on your server. For example, if you have a Hauppauge MediaMVP set-top box, you can run a special bootstrap load on it that will communicate directly with a MythTV back end on your server and let you watch your recorded content anywhere in the house.

Most users, however, also will want to run mythfrontend, which provides all the PVR user-level functionality through a GUI. In addition to letting the viewer choose what to record and to watch existing programming (as well as live shows), mythfrontend also can display weather data, current news, browse Web pages and even play games through the use of plugins. You can use multiple front ends (conceivably running on different machines), all talking to the same back-end server.

To begin with, we need to talk a bit about hardware. As with most things, what you are going to need depends on what you want to do with it. For example, the more tasks you expect your MythTV server to handle simultaneously, the more processor power you will need. Recording two shows at the same time while watching a third and transcoding a fourth for a DVD burn can take a fair amount of horsepower, so it doesn't hurt to spend a little for a decent processor. Thankfully, you won't need a liquid-nitrogen cooled, triple-overclocked speed demon to get the job done, a 2.8GHz chip should do it nicely.

You're also going to want a good supply of disk space on hand to store all those "Survivor: Sunnyvale" episodes. Because the whine of a noisy disk is the last thing you want to hear while watching your favorite show, go with SATA. A pair of 250GB drives shouldn't set you back more than $250 US if you wait for them to be on sale, and they will hold enough content to satisfy even the most hardened video junkie. We'll talk about filesystems in a bit.

Surprisingly, the video adapter is not a critical component. This doesn't mean you should dig out that old circa-1995 Hercules card you've got lying in the back room, but any reasonably recent AGP card of the past year should do just fine, although good OpenGL support will help out a lot. Tuner cards are a big deal, however. Choosing the right cards can make setup and use of your MythTV system much easier. You obviously need to look at factors such as whether you want to record HDTV, in which case picking a non-HD card is a non-starter. For the purposes of this article, we use the workhorse of most MythTV systems, the Hauppauge WinTV-PVR-250 and WinTV-PVR-350. What makes the Hauppauge cards so attractive is that they include MPEG encoders on the card, which drastically reduces the workload on the host CPU. You can easily record two shows at once using two of these cards and see only 5–6% CPU usage. The difference between the 250 and 350 is that the 350 also includes a hardware MPEG decoder and video out connector, so that you can hook it up to a TV set. However, it's going to run you another $50 US or so more than the 250. Because you probably don't want to stick your server in the middle of the living room just so you can hook it to your TV, I'd recommend going with the 250, which can be had for around $130 US retail and get a MediaMVP (around $80 US) for your TV hookup. Hauppauge also offers a WinTV-PVR-500 MCE with two tuners built in and a video out, but it doesn't include a remote, which is useful for controlling MythTV from a distance. The WinTV-PVR-250 is a rea-

Figure 1. The Hauppauge WinTV PVR-250 is the MythTV card of choice for analog TV reception.

sonably economical way to get shows onto your system, but be warned that if you're going to try to record HDTV, the Hauppauges aren't going to do the job for you.

Another thing you need to think about is whether you're going to need to control a satellite receiver or cable box to change channels. To make this work, you'll probably need an IR Blaster—a device that hooks up to your serial port and sends the proper commands to your set-top box. Also be aware that you can tune only one channel per set-top box, so if you want to record two shows at once, you're going to end up forking out for two boxes. This is the one big advantage that the DVR solutions offered by the cable and satellite companies have; they are built in to the set-top box, so this isn't an issue.

With the hardware requirements out of the way, it's time to provision the system. In spite of its reputation as a hard install, I've found that Gentoo offers the easiest overall experience in setting up MythTV. Use any of the standard tutorials to get a base Gentoo system up. The main thing you need do is make sure you set your filesystems up right. Assuming you bought two 250GB SATA drives, you really want to use the Logical Volume Manager (LVM) to turn most of the space into one large partition. I recommend doing the install normally but leaving most of drive 0 and all of drive 1 unassigned. So, you might use 10GB of drive 0 to set up your root, usr, swap and var space, leaving 240 remaining. Once your system is up and running, set up a 490GB LVM partition out of the remaining space.

The generally accepted wisdom is to use JSF (the IBM Journaling File System) as the filesystem for the partition on which you'll store your shows. This is because it offers the best performance when deleting large files—an activity that MythTV does frequently. This means you should make sure to compile JSF into your kernel (and not as a module). You also need to bake in LVM support.

The Gentoo Wiki site (see the on-line Resources) offers an excellent walk-through on setting up the kernel correctly to handle the integrated Hauppauge remote control and install the required packages. Once you have everything up and operational, you're ready to configure MythTV itself. Thankfully, the

ASA COMPUTERS
www.asacomputers.com
1-800-REAL-PCS

Hardware Systems For The Open Source Community–Since 1989
(Linux, FreeBSD, NetBSD, OpenBSD, Solaris, MS, etc.)

The AMD Opteron™ processors deliver high-performance, scalable server solutions for the most advanced applications. Run both 32- and 64-bit applications simultaneously

AMD Opteron™ Value Server—$795
• 1U 14.3" Deep
• AMD Opteron™ 240
• 512MB RAM Max 8GB
• 40GB IDE HDD
• 2x 10/100/1000 NIC
• Options: CD, FD or 2nd HD, RAID

Front I/O Dual AMD Opteron™ Cluster Node—$1,600
• 1U Dual AMD Opteron™ Capable Font I/O
• Single 240 AMD Opteron™
• 1GB RAM Max RAM 16GB
• 80GB HDD
• Dual PCI Expansion Slot

8 Hot Swap Bays in 2U AMD Opteron™—$1,950
• 1 of 2 AMD Opteron™ 240
• 512MB RAM Max 16GB
• 3x80GB IDE RAID # 5
• 2xGigE, CD+FD
• Options: SATA/SCSI, Redundant PS

No Frills AMD Opteron™ Storage Server—$8,450
• 6TB+ IDE/SATA Storage in 5U
• Dual AMD Opteron™ 240
• 512MB RAM
• 6TB IDE Storage
• Dual GigE, CD
• Options: SATA HDD, DVD+RW etc.

Your Custom Appliance Solution
Let us know your needs, we will get you a solution

"Your Logo Here"

Custom Server, Storage, Cluster, etc. Solutions
Please contact us for all type of SCSI to SCSI, Fibre to SATA, SAN Storage Solutions and other hardware needs.

AMD 64 Opteron

2354 Calle Del Mundo, Santa Clara, CA 95054
www.asacomputers.com
Email: sales@asacomputers.com
P: 1-800-REAL-PCS | FAX: 408-654-2910

Prices and availability subject to change without notice. Not responsible for typographical errors. All brand names and logos are trademark of their respective companies.

setup pretty much consists of running `mythsetup` and walking through a series of wizard screens that configure things such as your home cable/satellite system information. It shouldn't take more than five to ten minutes to do the basic setup.

One of the interesting things about MythTV is that it stores everything (except the actual video, of course) in a MySQL database. This makes it easy to import and extract information, and make tweaks. For example, if you need to fine-tune a channel's frequency, you can poke different values into the appropriate database table, go up and down one channel using the remote and see if it made things any better. Of course, it would be really snazzy to be able to tweak the fine-tuning using the GUI; maybe someone will implement that for a future release. It also means that you can run a simple query and see every show you've ever watched, or even write custom software that leverages the two-week program guide data MythTV automatically downloads for you.

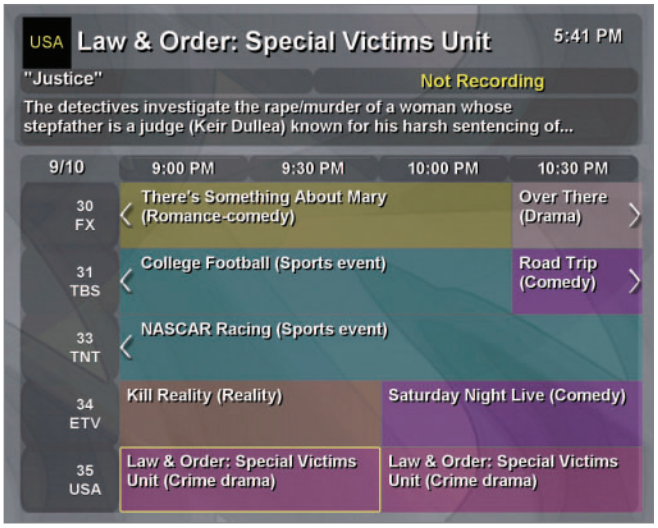

Figure 2. Program Guide

Once this is all working, you should be able to record shows and watch them on your monitor and speakers (which would be plugged in to your sound card and or motherboard speaker jack). To use the Hauppauge MediaMVP, you need to enable NFS on the server and export the filesystem with the video content. You also have to run a DHCP server and tftp server. Again, there's an excellent walk-through at the SourceForge site (see Resources).

Another option is to install MythWeb, which gives you an Apache-driven Web front end to view your program guide, scheduled recordings and already recorded programs. On Gentoo, this is as simple as typing `emerge mythweb`.

One outstanding feature of MythTV is the ability to skip commercials automatically. You enable this with a check box in the setup wizard. Once turned on, programs are queued up for commercial scanning after the end of the show. This means you can't skip commercials while watching a show that's being recorded or soon after, but generally the flagging is available within 30 minutes from the time the show ended. Then, while watching the show, you can use the skip-forward button to move past a group of commercials. The flagging isn't perfect, but it's pretty close. You also can set up MythTV to transcode content for later DVD burning automatically. In

Figure 3. MythWeb

Figure 4. Recorded Programs

fact, you can configure it to run any arbitrary Linux program on a video file after the recording is complete.

Is MythTV the right solution for you? If all you're looking to do is record content on the TV to which your PVR is attached, probably not. You can purchase DVRs from your cable or satellite providers that are cheaper and better integrated with their content and don't take a day or more to set up. But if having full control over your content is important, if you want to be able to share it all over your house from a single source and if you don't want to fork out $150 US a year for program data, MythTV offers the ultimate in flexibility, configurability and hackability.

Resources for this article: www.linuxjournal.com/article/8584.

James Turner is Product Review Editor for *Linux Journal*. He has written two books on Open Source Java development and is a Senior Software Engineer with Axis Technology, LLC.

LISA '05

19TH LARGE INSTALLATION SYSTEM ADMINISTRATION CONFERENCE

DEC. 4–9, 2005

SAN DIEGO, CA

www.usenix.org/lisa05/lja

FIND THE MISSING PIECES
TO YOUR TOUGHEST PUZZLES.

6 DAYS OF TRAINING
by industry experts, including:

- Rik Farrow on Hands-on Linux Security
- Don Bailey on 802.11 Wireless Network Penetration Testing
- Richard Bejtlich on Network Incident Response
- Jacob Farmer on Disk-to-Disk Backup

LISA '05 offers the most in-depth, real-world system administration training available!

3-DAY TECHNICAL PROGRAM

KEYNOTE: Qi Lu, Vice President of Engineering, Yahoo! Inc., on "Scaling Search Beyond the Public Web"

20+ INVITED TALKS INCLUDING:

- Matt Blaze, University of Pennsylvania: "Picking Locks with Cryptology"
- Kevin Bankston, EFF: "How Sysadmins Can Protect Free Speech and Privacy on the Electronic Frontier"

NEW! Hit the Ground Running Track

Refereed Papers, Guru Is In Sessions, Vendor Exhibition, BoFs, WiPs, and more!

Real World System Administration

SPONSORED BY

USENIX THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

SAGE The People Who Make IT Work

Register by November 18 and save!
www.usenix.org/lisa05/lja

VMware Workstation, Version 5.5

Virtualizers, rejoice! VMware has released version 5.5 of their popular VMware Workstation product. The new 64-bit guest support means that you can run either 32-bit or 64-bit versions of Linux, FreeBSD or Windows on a 64-bit host system, and additional import support lets developers load Symantec Ghost images as VMware virtual machines. The cost is $189 US for electronic download.

www.vmware.com

WinSystems' EPX-GX Single-Board Computer

WinSystems' new EPX-GX single-board computer offers a diversity of interface options for developers working on machine-to-machine applications. Based on an AMD GX500@1W processor, it draws a miserly 1.8 A at five volts, but still manages to sport a cornucopia of options, including 10/100 Ethernet, 802.11 support via a miniPCI connector, two USB ports, four comm ports, 24 digital I/O lines, audio, 4x AGP video, keyboard and mouse. Intended for use in applications such as robotics, transportation and other uses requiring a lower-power embedded device, it is compatible with the EPIC standard and is available for $499 US in OEM quantities.

www.winsystems.com

SUSE Linux, Version 10

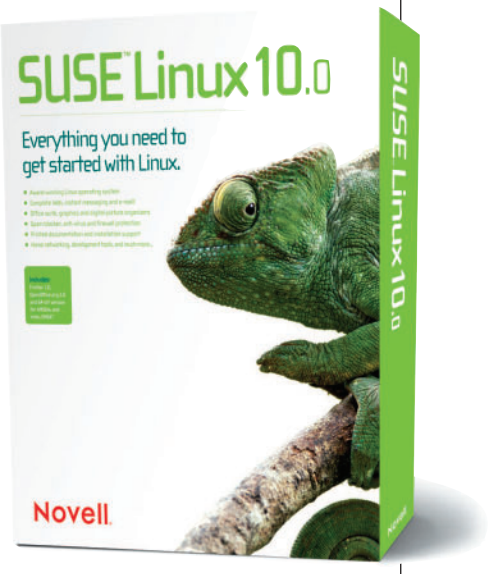

Novell has taken the wraps off the latest version of SUSE Linux. Version 10 is the first release to take advantage of Novell's new OpenSUSE initiative. OpenSUSE is Novell's answer to Fedora, letting community members contribute features and fixes to the SUSE Linux offering. Version 10 includes the latest versions of Firefox, OpenOffice.org and improved Windows integration, as well
as new features such as Xen virtualization and iFolder. Available for $59 US.

www.novell.com/products/linuxprofessional/resellers/index.html

Please send information about releases of Linux-related products to newproducts@ssc.com or New Products c/o *Linux Journal*, PO Box 55549, Seattle, WA 98155-0549. Submissions are edited for length and content.

Advanced MythTV Video Processing

Advanced methods for deinterlacing video playback and extracting video to take on the road.

BY MATTHEW GAST

Because of its stability and extensibility, Linux is often found at the frontier of computing. Linux has emerged as a promising platform for home theater audio/visual applications. My television viewing is now handled by a special-purpose Linux PC running MythTV. As I set up MythTV, the two major pain points I encountered were both related to video processing. The first challenge is to configure smooth video playback, and the second challenge is to take recorded programs on the road.

Deinterlacing Video Playback

To work within the limitations of the electronics of the day, television frames are transmitted as two separate "fields". A field consists of either the even-numbered or the odd-numbered horizontal lines in the picture. On playback, the even-numbered and odd-numbered fields are weaved together, and viewers far enough away from the display see continuous blended motion.

Two consecutive fields are related, but are not identical. During periods of rapid side-to-side motion of the camera, a field will be slightly ahead of its predecessor, and there may be jagged edges to images sliding across the screen. Figure 1 is a screenshot from a 1080i high-definition broadcast. In the scene, the camera is panning from left to right, causing the objects in the image to slide rapidly across the screen. Each field is in a slightly different position, leading to sawtooth-edge distortion, which is also called combing, serrations or mice teeth. In scenes with a great deal of sideways motion, it may be extremely difficult to follow the content through the distortion.

To make a video like Figure 1 watchable, it can be converted into a smooth picture by a process called deinterlacing. MythTV offers users a choice between several deinterlacing methods:

- One field—instead of using two fields for one frame, this extremely simple method keeps only one of the two fields. Every other field is displayed as a still image, and the unused fields are discarded.
- Linear—this method blends directly adjacent lines, which by definition come from alternate fields, together. A slight ghost image may appear, but the sawtooth distortion will be gone.
- Kernel—this method blends several lines together instead of just adjacent lines. Ghost images do not appear, though faint remnants of one field may remain.
- Bob—Bob is the most taxing method. Each field is line-doubled to create a frame, and then the reconstructed frames are played at double the frame rate.

Figure 1. Combing Distortion on Playback

Deinterlacing does require significant processing power, but most modern CPUs have multimedia instruction sets that reduce the load of the processing power. If you have an Intel processor with MMX or SSE instructions, or an AMD processor with 3Dnow!, deinterlacing should not be too difficult.

Bob is the best deinterlacing method to use with a synchronous TV output, though it can tax a less-capable machine. My personal MythTV front end is a 2GHz AMD Athlon64, and it has more than enough power to display Bob-deinterlaced high-definition video to an analog TV set. Although the CPU requirement is higher than the other deinterlacing methods, it is still well within the performance capabilities of my system.

Linear deinterlacing and kernel deinterlacing have similar visual effects, with the latter having a slightly larger CPU impact. Both are less taxing than Bob, which may be helpful on underpowered CPUs. Between the two, I prefer kernel deinterlacing because it blends several adjacent lines and eliminates ghosts, which make the resulting picture sharper.

Exporting Video to Other Systems

One of the initial reasons for setting up a MythTV system was a desire to take my television programs on the road. Now that

"television" means "files on hard disk", it is much easier to store, transport and watch where it it is convenient for me. Like many mobile professionals, my laptop has practically become an extra appendage, and it is an ideal platform for playing back video while mobile, especially now that many airplanes have added power ports for mobile electronics.

The basic tension in exporting video from MythTV is a trade-off between size and processing time. Digital TV broadcast standards describe how to send an MPEG-2 video stream over a TV channel, so it is natural for MythTV to store digital TV broadcasts in their "natural" MPEG-2 format. Converting the MPEG-2 digital TV stream to another MPEG-2-based video format is relatively easy and can be done without lots of processing time. Converting the digital TV stream to MPEG-4 requires much more processing power, but the resulting video file will be much smaller.

Table 1. Results of Export (Typical)

Format	Resolution	File size	Encoding time
Myth native format	704x480	1,756MB	N/A
Commercial DVR	N/A	1,236MB	N/A
VCD	352x240	596MB	15 min.
SVCD	480x480	601MB	25 min.
DVD	720x480	899MB	34 min.
DivX	624x464	432MB	35 min.
XviD (1 pass)	624x464	451MB	1 hour, 39 min.
XviD (2 pass)	624x464	472MB	2 hours, 35 min.
ASF	320x240	143MB	18 min.

Extracting the MythTV Recordings

Although digital TV recordings are an MPEG-2 video stream, the NuppelVideo container format used by MythTV is specific to MythTV and is not supported by most video player software. To watch the videos with anything other than a MythTV front end, you must convert them to a format with a wider selection of players.

Exporting video is further complicated by the filenames used by MythTV to store recordings. The first part of the filename is the channel number used by MythTV, and the two long numbers are the start date and end date:

```
myth:/video$ ls -lh *.nuv
-rw-r--r--  1 mythtv users 2.8G Sep  3 08:00 1000_20050903073000_20050903080000.nuv
-rw-r--r--  1 mythtv users 9.1G Aug 31 23:30 1007_20050831220000_20050831233000.nuv
-rw-r--r--  1 mythtv users 808M Sep  3 01:30 1008_20050903010000_20050903013000.nuv
-rw-r--r--  1 mythtv users 1.8G Sep  1 09:00 1009_20050901080000_20050901090000.nuv
-rw-r--r--  1 mythtv users 3.7G Aug 28 22:00 1019_20050828210000_20050828220000.nuv
```

Tools to convert the MythTV recordings into more widely supported formats are readily available. mythtranscode, a program that can decode the MythTV NuppelVideo files into a standard video stream supported by many programs, comes as part of the MythTV distribution. By feeding the output of mythtranscode to an encoder program of your choice, you can create any type of video file you want. A common choice of encoder program is FFmpeg, which supports many common video formats. Linking the two programs together to produce an intelligible output file is theoretically possible, but wading through all the command-line options and the recordings on your system is a complicated process.

Enter nuvexport, a Perl script that manages the process for you in an extremely friendly manner. nuvexport assembles the command-line options necessary to run mythtranscode and FFmpeg. It uses a text-based menu interface to guide you through the process of selecting a show for export and then setting up any required parameters of the conversion programs.

Before installing nuvexport, there are a few necessary support tools. FFmpeg is the default program required to re-encode the video into the selected target format. (nuvexport also supports transcode, which is much slower.) MPlayer is used to decode the MythTV files for conversion. Many conversions result in the appearance of noise that needs to be filtered

out. nuvexport uses the yuvdenoise program, which is one of the MJPEG tools. All three programs are widely used and are likely to be available as packages for your Linux distribution. nuvexport uses the DateManip module as well, so fetch that from your distribution's package site or favorite CPAN mirror.

The first prompt from nuvexport is to select the format of the exported video file. The basic trade-off is whether to choose an MPEG-4-based export format for minimum file size at the cost of extra processing time to prepare it, or to use one of the larger but easier to prepare formats. The major choices are:

- Video CD (VCD) consists of an MPEG-1 video stream at 1,150kbps, with the audio in a separate MPEG layer 2 (MP2) track at 224kbps.
- Super Video CD (SVCD) consists of an MPEG-2 video at a variable rate, while retaining the MP2 audio track. Unlike VCD, the audio track can have multiple channels, so 5.1 or 7.1 audio can be stored in this format.
- DVD is based on MPEG-2 video, with several options for an audio track. It is higher resolution than either of the video CD formats.
- DivX is an MPEG-4-based format, which results in small files. However, it minimizes file size without significant sacrifices in quality. DivX can produce either constant-rate video or variable-rate video.
- XviD is an MPEG-4-based format that is an offshoot of the development of DivX. It is based on an open-source development of a DivX codec released in 2001. By default, nuvexport uses variable-rate video encoding for XviD and offers the option of either a single pass or multiple passes. Multiple passes improve video quality at the expense of processing time.
- Advanced Streaming Format (ASF) was developed by Microsoft as a generic container for media, and it is commonly used with Windows Media Audio (WMA) and Windows Media Video (WMV) files. WMV is based on Microsoft extensions to MPEG-4.

In my experience, the VCD and SVCD codecs offer good quality with fast processing times, while the DivX and XviD codecs offer the smallest file size but take longer to produce. After selecting a video export format, nuvexport uses a text-based menu system to select the episodes for export and set up parameters for the codec.

After selecting a set of episodes for export, nuvexport presents standard questions, such as where to put the exported file. It offers the option of using the MythTV cutlist, which cuts commercials from the exported video. Noise reduction and deinterlacing are offered as options. Although both default to yes, I usually disable them because of the additional processing time. Many video players can deinterlace on playback, and I have not found excessive noise from conversion.

Some export formats have additional codec-specific questions. Both DivX and XviD will allow adjustment of the bit rates and resolution. The default bit rates of 128kbps for audio and 960kbps for video are sufficient to produce good-quality

Systems Management: Clusters and Supercomputer for Computational Biochemistry

Extraordinarily gifted individuals sought to provide Linux systems administration and networking support for a rapidly growing New York–based technology project aimed at achieving major scientific advances in the field of biochemistry and fundamentally transforming the process of drug discovery. This research effort is being financed by the D. E. Shaw group, an investment and technology development firm with approximately $17 billion in aggregate capital, and operates under the direct scientific leadership of its founder, Dr. David E. Shaw.

Successful hires will be responsible for operational support for and substantial research projects within our Linux and network (Cisco, Infiniband) environments, including one of the largest Linux clusters in the world, as well as a massively parallel specialized supercomputer incorporating 90-nanometer "system on a chip" ASICs. Ideal candidates will have a computer science, engineering, or science degree, extensive knowledge of multiple Linux/UNIX operating systems, strong programming and scripting ability, and excellent verbal and written skills. We are prepared to reward exceptionally well-qualified individuals with above-market compensation.

Please send your resume to linuxjournal-sa@desrad.deshaw.com.

Members of the D. E. Shaw group do not discriminate in employment matters on the basis of race, color, religion, gender, national origin, age, military service eligibility, veteran status, sexual orientation, marital status, disability, or any other protected class.

D E Shaw & Co

video on most computer displays. When setting video size, nuvexport prompts for the width first and then proposes a height based on the aspect ratio of the recording. Keep the width less than the recorded width. The default width of 624 usually produces good video, but it can be larger for recordings that are 1920x1080. VCD and SVCD do not prompt for resolution because the formats have fixed resolutions.

Table 1 compares the processing time and space required by each of the different formats, as well as the MythTV native format and my commercial DVR. As a source, I used an hour-long episode of PBS's *Nova*, which is transmitted at a resolution of 704x480. The "commercial DVR" entry is for the video program as transferred from my commercial DVR to my laptop. As a rule of thumb, MPEG-2 requires approximately 1GB per hour, but MPEG-4 will be only 350–450MB. The encoding time listed in the table does not take into account either deinterlacing or noise reduction; I perform both tasks in the video player on playback.

Although ASF has the smallest size, it is also by far the worst looking. There are large compression artifacts in the ASF file that make it very distracting to watch. Although the small size is attractive, the poor picture quality rules it out. DVD video has the best picture quality, but it also requires the most disk space. As a compromise between the two extremes, I use VCD and DivX, depending on priorities. I use the former to create a file quickly, and the latter to create the smallest file possible.

Video transcoding is a CPU-intensive process. By default, nuvexport runs its helper processes at high nice values to prevent them from interfering with other system operations, such as video playback or recording. All recent Linux distributions have software that allows the CPU clock speed to be changed in response to demand for processing power. I use the CPU speed control to keep the clock speed as low as possible while still accomplishing the work I want. Many CPU speed control programs will not take into account niced processes, but they can be configured to do so. My Linux distribution uses the CPUfreq kernel driver, which needs to be configured to monitor niced processes. A small start-up script runs the following two commands:

```
echo "ondemand" > /sys/devices/system/cpu/cpu0/cpufreq/scaling_governor
echo "1" > /sys/devices/system/cpu/cpu0/cpufreq/ondemand/ignore_nice
```

The second command instructs CPUfreq driver to count the processing demands from niced processes. Times listed in Table 1 are from a 2.0GHz AMD Athlon64 running at top speed. At the minimum speed of 1.0GHz, processing time is approximately three to four times as long. Exporting the video from MythTV is only the first half of the battle. Once the video is produced, it can be transferred to another location for playback. In addition to the playback applications present on the viewing platform, there are two notable open-source playback applications: MPlayer and Video LAN Client (VLC). I use MPlayer because the built-in deinterlacing capabilities result in a smoother picture than VLC. Both applications are available on both Linux and Windows.

MPlayer's command line is identical on different host operating systems. The goal is to get crisp full-screen video playback. The -fs option plays back the video with the full screen so there is no window around the video. Video filters can be used to change the playback and are activated with the -vf option. I use two -vf options. One creates a small black border around the screen with the expand filter. The expand filter takes multiple arguments. A negative number is interpreted as a border. The filter `-vf expand=0:-50` puts a 50-pixel border at the bottom of the screen and leaves the video centered in the border. To get crisp video, deinterlacing is necessary. MPlayer activates deinterlacing with the postprocessing filter, abbreviated pp. As a general rule, I turn on four postprocessing filters: horizontal de-blocking (hb), vertical de-blocking (vb), de-ringing (dr) and brightness/contrast correction (al). The resulting filter is activated with `-vf pp=hb/vb/dr/al`. Putting it all together, the command line is:

```
mplayer -fs -vf pp=hb/vb/dr/al -vf expand=0:-50 (filename)
```

During playback, MPlayer's extensive keyboard commands enable pausing, fast-forwarding and picture adjustment, as well as on-screen display.

Resources for this article: www.linuxjournal.com/article/8585.

Matthew Gast is the author of the leading technical book on wireless LANs, *802.11 Wireless Networks: The Definitive Guide* (O'Reilly Media). He is currently spending far too much time these days working with MythTV. He can be reached at matthew.gast@gmail.com, but only when he is close to sea level.

Monitoring Virtual Memory with vmstat

Just using a lot of swap space doesn't necessarily mean that you need more memory. Here's how to tell when Linux is happy with the available memory and when it needs more. BY BRIAN K. TANAKA

Linux novices often find virtual memory mysterious, but with a grasp of the fundamental concepts, it's easy to understand. With this knowledge, you can monitor your system's memory utilization using vmstat and detect problems that can adversely affect system performance.

How Virtual Memory Works

Physical memory—the actual RAM installed—is a finite resource on any system. The Linux memory handler manages the allocation of that limited resource by freeing portions of physical memory when possible.

All processes use memory, of course, but each process doesn't need all its allocated memory all the time. Taking advantage of this fact, the kernel frees up physical memory by writing some or all of a process' memory to disk until it's needed again.

The kernel uses paging and swapping to perform this memory management. Paging refers to writing portions, termed pages, of a process' memory to disk. Swapping, strictly speaking, refers to writing the entire process, not just part, to disk. In Linux, true swapping is exceedingly rare, but the terms paging and swapping often are used interchangeably.

When pages are written to disk, the event is called a page-out, and when pages are returned to physical memory, the event is called a page-in. A page fault occurs when the kernel needs a page, finds it doesn't exist in physical memory because it has been paged-out, and re-reads it in from disk.

Page-ins are common, normal and are not a cause for concern. For example, when an application first starts up, its executable image and data are paged-in. This is normal behavior.

Page-outs, however, can be a sign of trouble. When the kernel detects that memory is running low, it attempts to free up memory by paging out. Though this may happen briefly from time to time, if page-outs are plentiful and constant, the kernel can reach

a point where it's actually spending more time managing paging activity than running the applications, and system performance suffers. This woeful state is referred to as thrashing.

Using swap space is not inherently bad. Rather, it's intense paging *activity* that's problematic. For instance, if your most-memory-intensive application is idle, it's fine for portions of it to be set aside when another large job is active. Memory pages belonging to an idle application are better set aside so the kernel can use physical memory for disk buffering.

Using vmstat

vmstat, as its name suggests, reports virtual memory statistics. It shows how much virtual memory there is, how much is free and paging activity. Most important, you can observe page-ins and page-outs as they happen. This is extremely useful.

To monitor the virtual memory activity on your system, it's best to use vmstat with a delay. A delay is the number of seconds between updates. If you don't supply a delay, vmstat reports the averages since the last boot and quit. Five seconds is the recommended delay interval.

To run vmstat with a five-second delay, type:

```
vmstat 5
```

You also can specify a count, which indicates how many updates you want to see before vmstat quits. If you don't specify a count, the count defaults to infinity, but you can stop output with Ctrl-C.

To run vmstat with ten updates, five seconds apart, type:

```
vmstat 5 10
```

Here's an example of a system free of paging activity:

```
procs                  memory    swap        io     system cpu
r  b  w   swpd   free  buff  cache  si so   bi  bo   in    cs us sy  id
0  0  0  29232 116972  4524 244900   0  0    0   0    0     0  0  0   0
0  0  0  29232 116972  4524 244900   0  0    0   0 2560     6  0  1  99
0  0  0  29232 116972  4524 244900   0  0    0   0 2574    10  0  2  98
```

All fields are explained in the vmstat man page, but the most important columns for this article are free, si and so. The free column shows the amount of free memory, si shows page-ins and so shows page-outs. In this example, the so column is zero consistently, indicating there are no page-outs.

The abbreviations so and si are used instead of the more accurate po and pi for historical reasons.

Here's an example of a system with paging activity:

```
procs                  memory    swap        io     system cpu
r  b  w   swpd   free  buff  cache  si  so   bi  bo   in    cs us  sy  id
. . .
1  0  0  13344   1444  1308  19692   0 168  129  42 1505   713 20  11  69
1  0  0  13856   1640  1308  18524  64 516  379 129 4341   646 24  34  42
3  0  0  13856   1084  1308  18316  56  64   14   0  320  1022 84   9   8
```

Notice the nonzero so values indicating there is not enough physical memory and the kernel is paging out. You can use top and ps to identify the processes that are using the most memory.

You also can use top to show memory and swap statistics. Here is an example of the uppermost portion of a typical top report:

```
14:23:19 up 348 days,  3:02,  1 user,  load average: 0.00, 0.00, 0.00
55 processes: 54 sleeping, 1 running, 0 zombie, 0 stopped
CPU states:   0.0% user,   2.4% system,   0.0% nice,  97.6% idle
Mem:    481076K total,   367508K used,   113568K free,     4712K buffers
Swap:  1004052K total,    29852K used,   974200K free,   244396K cached
```

For more information about top, see the top man page.

Conclusion

It isn't necessarily bad for your system to be using some of its swap space. But if you discover your system is often running low on physical memory and paging is causing performance to suffer, add more memory. If you can't add more memory, run memory-intensive jobs at different times of the day, avoid running nonessential jobs when memory demand is high or distribute jobs across multiple systems if possible.

Resources for this article: www.linuxjournal.com/article/8535.

Brian K. Tanaka has been a UNIX system administrator since 1994 with companies such as SGI, Intuit and RealNetworks. He is cofounder of Martingale-Oak, LLC. He can be contacted at btanaka@martingale-oak.com.

LAYER 42 TM
Redundant UPS and generator
Nationwide network
Free tech support
2U
256kbps ~80GB
$60/mo.
4U or Mid-tower
256kbps ~80GB
$80/mo.
1/4 Rack
512kbps (14U) ~165GB
$200/mo.
1/2 Rack
1mbps (28U) ~330GB
$350/mo.
www.layer42.net
All prices include 100Mbps port, Firewall,
24x7 Monitoring and DNS hosting
408-450-5740 2336-F Walsh Ave., Santa Clara, CA 95051

Making Linux Accessible for the Visually Impaired with Speakup

Speakup makes Linux more accessible to the visually impaired by integrating speech capabilities directly into the kernel.

BY AMEER ARMALY

During the past ten years, evolutions in many fields of technology have influenced the lives of all of us, and especially the world's blind population. Advancements in speech synthesis have led to the usability of many different operating systems, Linux among them. One of these programs, and by far one of the best, is a screen review package called Speakup, written by Kirk Reiser with assistance from the user community. Speakup is unique in the sense that it integrates seamlessly into the kernel, allowing it to talk from startup to shutdown, and even to debug kernel errors, which I can testify to from personal experience. It also makes the installation of a Linux system much easier, because one does not usually require a serial console or sighted assistance to complete the installation process.

A screen review package is a program that takes the text displayed on the screen, and outputs it in spoken words. The actual speaking is done by a speech synthesizer, which can come in either hardware or software versions. Hardware synthesizers are either external boxes with headphone jacks and volume knobs that plug in to your computer via serial or USB ports, or ISA or PCI cards that have an output jack for a speaker or headphones. Software synthesizers are actual software programs that handle all the processing of the text into spoken words and output it through the computer's sound card. Speakup supports both hardware and software synthesizers, though software synthesizers require a user-space program and thus can't load at kernel boot, as we'll discuss later. Speakup's key features include seamless integration, logical key layout, support for laptop keyboards, easy adjustability of speech settings and support for software synthesizers.

Features

Speakup is packed full of features, some of which you won't find in any other screen reader. In order to read text, Speakup uses an invisible review cursor. At the same time, however, Speakup tracks the system cursor, to facilitate navigation in menus, editors and similar situations. To perform tasks such as moving the review cursor around, Speakup uses the numeric keypad, hereafter referred to as the numpad.

The numpad Enter key silences speech until the next key press, which is very useful for quieting boot-up messages and/or frequently heard text. It also synchronizes the location of the review cursor with the system cursor, facilitating many different operations. Insert plus numpad Enter silences reading of new text until this combination is pressed again, but still allows you to move around the screen.

The numpad plus key reads the

NCBI
EMBOSS
Glimm
Clust W
HPC
Bioinformatics Package
HMMER
Wise
FastA
Enhanced Pise foundation.
Elegant web interface.
Integrated file manager.
Secure Data Access Management.
Zero post-installation configuration.
Rapid cluster construction.
Simple node image updates.
Optimized for 64-bit Power Linux.
Account Administration.
Value Added Reseller
IBM Business Partner
www.terrasoftsolutions.com

entire screen. The numpad 0, or insert, is used as a key modifier similar to Alt, Ctrl or Shift. Speakup also respects numlock, still allowing the user to enter numbers from the numpad if necessary. Numpad keys 7–9 go up a line, read the current line and go down a line, respectively. A similar arrangement is used for words on numpad 4–6, and with characters on numpad 1–3. The numpad slash marks a spot on the screen, and if there is a spot already marked, it copies the text into memory. Insert plus numpad slash inputs any previously copied text, which usually results in pasting it to the location of the system cursor.

The numpad minus parks the review cursor. Parking means that the review cursor's location will not be moved unless the user moves it; this is useful for tracking text that changes but is not at the cursor, requiring you to move to it constantly. This functionality is also in the windowing system, which will be covered shortly. Numpad star toggles on and off cursor tracking. This is different from parking the review cursor, because parking does not affect what is actually spoken, just where the review cursor is. Cursor tracking always speaks what is at the cursor, which is optimum for menus and editors, but occasionally you may need to turn it off.

Laptops

For laptops, Speakup has a set of key assignments as well. These center around the Caps Lock key or Windows logo key if it is present on the keyboard. While the Caps Lock key is down, the letters I, O and U act as the numpad 7–9. Thus, you have a very similar arrangement to what you have on the numpad. Some things are different—for instance, Caps Lock plus Enter acts as numpad Enter, but overall it's very similar and easy to learn. When referring to either the the Caps Lock/Windows key or numpad Insert key simultaneously, they are called the Speakup key.

Adjusting Settings

Adjusting speech settings, such as volume, rate, pitch and tone, can be done in two ways.

The first, and probably the easiest, is to use the Speakup key plus the numbers on the number row. The Speakup key plus 1 and 2 decrement and increment the volume, respectively; 3 and 4 do the same with pitch; and finally 5 and 6 do the same with rate. The Speakup key plus F9 and F10 control punctuation, and the Speakup key plus F11 and F12 control the punctuation only for reading.

The Speakup key plus F5 lets you edit the "some" punctuation level. It works by toggling the punctuation that you press, as to whether it is spoken in the specified level. The Speakup key plus F6 does the same for the "most" punctuation level, and Speakup key plus F7 lets you edit what delimiters are used when moving by words; usually it is spacing and certain punctuation.

The other method of changing speech settings is to use the Speakup entry under /proc. Under /proc/speakup, there are the usual items, such as volume, rate, pitch, voice, version and synth_name, as well as some more-advanced items dealing with timing and other things. Some of these values are read/write, and some are read-only. For instance, version gives the current revision of Speakup, including the CVS build date if applicable, but synth_name can be used both to get and set the synthesizer in use. synth_direct is a write-only entry that sends all text directly to the synthesizer. It is even possible to load a new keymap while the system is running, rather than having to rebuild the kernel. There are also values for punct_some, punct_most and delimiters, which do the same things as the key functions described above. There is also a script called speakupconfig, which saves all of your entries in /proc/speakup for the particular synthesizer in use and allows you to restore these settings later, allowing automated loading of settings.

Windows

Speakup has a windowing system, which can be very useful in certain programs where a specific area of the screen that is not tracked by the cursor is updated frequently. The Speakup key plus F2 is used to set the window dimensions; the Speakup key plus F3 clears the window settings, allowing you to set a new one; and the Speakup key plus F4 silences the window, preventing it from being read automatically. However, you can read windows manually with the Speakup key plus the numpad plus key.

Work is now being done on color and highlighting recognition, which will allow ncurses-based programs to function even better than they do now, especially in menus. This means that text that is a different color from surrounding text will be given a higher priority, thus read first.

Help

There are several ways to get help on Speakup. First, you can load the module called speakup_keyhelp, and press the Speakup key plus F1. This puts you in a key identification mode, which can be exited by pressing the spacebar. When in this mode, Speakup speaks the description of any key that is assigned to a Speakup function, and allows you to arrow through the list of assignments. Another way to get help is to consult the guide provided with Speakup under Documentation in the kernel tree, or on the Web site. This document has many useful instructions, which can get a new user started with Speakup, as well as refresh an existing user's memory.

Installation

The number one thing that sets Speakup apart from other screen reader programs is the fact that it is literally part of the kernel. The install script applies a few patches to some kernel source files and copies the relevant Speakup sources to drivers/char in the kernel tree. Then, when `make config` is executed, there is a section for console speech output and Speakup. There you can choose what synthesizers you would like to build directly in to the kernel or as modules, though software speech support can be built only as a module.

You can also select what synthesizer you want to be the default at startup. Thus, if you build everything in to the kernel, you have a fully talking Linux system from startup to shutdown. This allows a blind person to install Linux without any sighted assistance whatsoever, because every step in the installation talks.

There are Speakup-modified ISO images for three major distros: Debian, Fedora and Slackware. Slackware has actually incorporated Speakup into its official installation setup, simplifying things even further. There is also a Speakup-enabled version of Knoppix, which is a basic Linux distro on CD. This allows people wanting a quick look at a Linux system simply to boot the CD, have it come up talking and not have to worry about installation unless

they're interested. It also can be very useful for crash recovery.

Software Speech

As previously mentioned, Speakup supports software speech synthesizers with some user-space support. Some of the more famous software synthesizers include Festival, Flite, Freetts and IBM's VivaVoice Outloud, which is no longer supported. Software speech in Speakup centers around another program called Speech Dispatcher. Speech Dispatcher is a framework to provide a single interface to multiple software synthesizers. It does this through a series of programs that provide a Speech Dispatcher interface to elements such as Emacs as well as libraries for a number of languages. It also has a tcp protocol for transmitting speech from a server to client that does the actual output.

Speakup has a generic software synthesizer driver called /dev/softsynth, which outputs the text that would normally be sent to a hardware synthesizer. A module for Speech Dispatcher, called speechd-up, takes the text from /dev/softsynth and sends it to Speech Dispatcher and a software synthesizer of the user's choice. Support exists for Festival, Flite, Dectalk software and generic synthesizers. You also can integrate other synthesizers with some tweaking of configuration files. Performance-wise, software synthesizers have a slight lag in responsiveness compared to hardware synthesizers, but the overall result is not that bad given the circumstances.

The first step is to get Speech Dispatcher working, which is not hard at all; just compile it and you're set to go. You have to edit the configuration file to tell it what synthesizer you want to use; by default it uses Flite. Then, compile and install speechd-up. To start software speech, load the speakup_sftsyn module if you haven't already, and run speechd-up. If you do this through an init script, you still will get an early-talking system, though not entirely in the kernel.

Future

Many things are planned for Speakup in the future. As has been previously mentioned, work has been started on color recognition and highlight tracking, thanks to some folks at the American Printing House for the Blind. This will enable many menu-based programs to talk much more smoothly.

Another new feature that is planned is keyboard macros, allowing the user to accomplish many different tasks with the press of one key. We eventually want to have a screen memory find function, as well as a goto function to go to a specific set of coordinates on the screen.

Another matter that is under consideration and analysis is configuration files. These files would somehow have to be loaded in on execution of their corresponding program, and would contain voice, macro and other information necessary for the operation of that program.

All of these and more features are planned for Speakup in the future, provided that people are willing to help and contribute their time to the effort of making Linux accessible to the world's blind population.

Conclusion

Today, technology has revolutionized the lives of the world's blind population. Computers allow us to access data more easily than ever, and the arrival of the Internet into the mainstream has made communication and linking with others easier than ever before for everyone. Linux systems are economical by their nature, not requiring the absolute latest hardware to run well. This is especially helpful for the world's blind, who may not have access to as much funding as would be ideal. Now there is a cheap and workable solution for those people, a fully talking Linux system with Speakup; and with the introduction of software speech and Speech Dispatcher, it just got even cheaper.

Resources for this article: www.linuxjournal.com/article/8586.

Ameer Armaly is a sixteen-year-old junior in high school. He has been blind since birth, and enjoys programming, food and science fiction. He uses computers with the aid of talking programs that read the text aloud, sometimes as fast at 550 words per minute.

UNIX: Old School

Experience historical UNIX releases firsthand using the SIMH simulator on Linux. BY MATTHEW HOSKINS

I have been called "nostalgic beyond my years" by some, and I suppose that is accurate. I was born in 1976 and have always had a voracious appetite for early minicomputer and mainframe history. I believe recorded history itself is the single-most important innovation of human existence. We humans seem to have a hard-wired compulsion to record, pass on and learn from the mistakes and successes of those before us. Open-source software is the natural evolution of this concept applied to computer technology. In the Open Source philosophy, we are all free to learn from the wealth of software created by the masses that came before us. By examining the evolution of a project, we can learn from the mistakes of others and, perhaps most important, copy verbatim from their successes. By harnessing this freely available history as well as unfettered cooperation, we advance the common good.

Recently, companies have begun to loosen their grip on their early computing "intellectual property". Although some have not fully embraced open source, these sometimes small, token gestures offer us a wealth of knowledge. In this article, I focus on how we can explore early operating system history by running "historic" UNIX releases on our very own Linux boxes using a simulator. The SCO Group (Yes, "them", previously Caldera, Inc.) claims current ownership of early UNIXes and has released them under an "Ancient Unix" license, which allows for noncommercial use. I focus here on the UNIX V5 release, because it is the earliest available. UNIX V6, V7 and various early BSD releases are also available. If you plan on trying out any of these OSes, examine the licenses included with each before booting them up.

In order to explore these OSes, we need to be able to run them on commonly available computing hardware. Luckily, we have simulators for this purpose. Because of its quality and depth of support, one of the most popular simulators is SIMH, available from the SIMH Web site (see the on-line Resources). SIMH runs on every popular *nix OS, as well as Microsoft Windows, and is capable of simulating a wide range of early computer systems, including Digital Equipment Corp.'s PDP and VAX systems, the MITS Altair, early IBM systems and many more. Some of the most historically significant systems are DEC's PDP series, the birth-system of UNIX.

SIMH is a ground-up system simulator; it simulates the

STRANGER IN A STRANGE LAND: THE UNIX V5 USER ENVIRONMENT

The UNIX V5 system provided in the disk image is rather stark and unfriendly compared to modern, lush UNIX/Linux systems. Here are a few pointers to get you started:

- sh is the shell. It's only 858 lines of C; don't expect it to work like bash.
- Use `chdir` to change the default directory.
- Backspace and arrow keys rarely work.
- ed is the text editor; see **en.wikipedia.org/wiki/Ed**.
- bas is a basic interpreter.
- fc is a FORTRAN interpreter.
- cc is the C compiler.
- Source code is in /usr/source.
- There are not many files, so use `find / -print` to see what else is included.

CPU, memory, firmware and devices of a number of early computer systems. This means that original distributed software can run unmodified on these simulated systems. SIMH successfully simulates devices such as disks, tape drives, printers and networking devices. This means that not only can we run these historic systems, but we can communicate and transfer data to and from them using modern technologies and protocols. A great deal of thanks is owed to the contributors of SIMH. Their decision to contribute and release under open source furthers all our understanding of our history and guarantees that this history will always be free.

Getting Started: Installing SIMH

Download the latest SIMH release, V3.4-0 at the time of this writing, compile and install. If you want to use Ethernet emulation, you may need to upgrade the libpcap library bundled with your OS as most currently distributed versions are too old. The SIMH installation documents explain how to do this, and you can skip this step if you're not going to be using networking support on your simulated machines. Compiling can be done as any user and is as simple as:

```
$ mkdir simh
$ cd simh
$ unzip /path/to/simhv34-0.zip
$ mkdir BIN  # Note all CAPS
$ gmake USE_NETWORK=1 all
# Only include USE_NETWORK=1 if your PCAP lib is up to
date.

(compilation chatter omitted)
$ ls -l ./BIN/
total 11624
-rwxrwxr-x  1 matt matt  301959 Jul 16 18:45 altair
-rwxrwxr-x  1 matt matt  482274 Jul 16 18:45 altairz80
-rwxrwxr-x  1 matt matt  529317 Jul 16 18:44 eclipse
-rwxrwxr-x  1 matt matt  297590 Jul 16 18:45 gri
-rwxrwxr-x  1 matt matt  375737 Jul 16 18:44 h316
-rwxrwxr-x  1 matt matt  577678 Jul 16 18:44 hp2100
-rwxrwxr-x  1 matt matt  355225 Jul 16 18:44 i1401
-rwxrwxr-x  1 matt matt  381672 Jul 16 18:45 i1620
-rwxrwxr-x  1 matt matt  441079 Jul 16 18:46 ibm1130
-rwxrwxr-x  1 matt matt  502037 Jul 16 18:46 id16
-rwxrwxr-x  1 matt matt  508378 Jul 16 18:46 id32
-rwxrwxr-x  1 matt matt  294614 Jul 16 18:46 lgp
-rwxrwxr-x  1 matt matt  434940 Jul 16 18:44 nova
-rwxrwxr-x  1 matt matt  345034 Jul 16 18:41 pdp1
-rwxrwxr-x  1 matt matt  752055 Jul 16 18:43 pdp10
-rwxrwxr-x  1 matt matt 1055376 Jul 16 18:43 pdp11
-rwxrwxr-x  1 matt matt  474153 Jul 16 18:42 pdp15
-rwxrwxr-x  1 matt matt  459203 Jul 16 18:41 pdp4
-rwxrwxr-x  1 matt matt  460363 Jul 16 18:41 pdp7
-rwxrwxr-x  1 matt matt  499473 Jul 16 18:42 pdp8
-rwxrwxr-x  1 matt matt  467662 Jul 16 18:42 pdp9
-rwxrwxr-x  1 matt matt  352233 Jul 16 18:45 s3
-rwxrwxr-x  1 matt matt  429312 Jul 16 18:46 sds
-rwxrwxr-x  1 matt matt  982694 Jul 16 18:43 vax
```

SUBSCRIBE
www.linuxjournal.com

SIMH is a ground-up system simulator; it simulates the CPU, memory, firmware and devices of a number of early computer systems.

This builds all possible system simulators. Each simulator becomes a separate binary in the ./BIN/ directory. SIMH can be run as any normal user, but if you want to use Ethernet network simulation, you need to execute it as root (under UNIX) to allow libpcap access to the Ethernet device.

Running UNIX V5

UNIX V5, released in June 1974, was still very early in UNIX development at Bell Labs. Much of the system was still written in assembler. This disk image includes a working C compiler (cc) and a great deal of interesting source code under /usr/source. To begin our exploration, we must download the UNIX V5 disk image (see Resources). This zip archive contains the pre-installed image file as well as a README and file containing license information. The disk image is a snapshot of a working installed system. In this case, it is simulating an RK05 disk drive. We must now collect the pieces we need to get this system booted. Begin by creating a directory, then copy the BIN/pdp11 binary from under the SIMH build directory as well as the contents of the uv5swre.zip archive uncompressed. Then, create a pdp11.ini file to control the simulator, using an editor of your choice, and place the following lines in the ini file:

```
set cpu U18
attach rk0 unix_v5_rk.dsk
boot rk0
```

This tells the simulator what kind of CPU to emulate and to attach the unix_v5_rk.dsk file as a simulated RK-style disk using the rk0 device name. Finally, this file tells the simulator to boot the OS image on that disk.

Your simulator directory should look like the following:

```
-rw-rw-r--  1 matt matt   12299 Jan 24  2002 AncientUnix.pdf
-rwxrwxr-x  1 matt matt  913614 Jul 22 19:33 pdp11
-rw-rw-r--  1 matt matt      47 Jul 22 23:59 pdp11.ini
-rw-rw-r--  1 matt matt     263 Nov 25  1996 README.txt
-rw-rw-r--  1 matt matt 2494464 Jul 23 00:39 unix_v5_rk.dsk
```

To boot up UNIX V5, simply type `./pdp11` in the current directory, then when prompted, type `unix` at the @ prompt. You almost immediately will get the login: prompt; there was not much in the way of boot messages in these old UNIXes. There is no root password, so you will be given a command prompt. Your session could look as follows:

```
$ ./pdp11

PDP-11 simulator V3.4-0
Disabling XQ
@unix

login: root

# ls -l /
total 60
drwxr-xr-x  2 bin        944 Nov 26 18:13 bin
drwxr-xr-x  2 bin         80 Nov 26 18:13 dev
drwxr-xr-x  2 bin        240 Mar 21 12:07 etc
drwxr-xr-x  2 bin        224 Nov 26 18:13 lib
drwxr-xr-x  2 bin         32 Nov 26 18:13 mnt
drwxrwxrwx  2 bin        112 Mar 21 12:11 tmp
-rwxrwxrwx  1 bin      25802 Mar 21 12:07 unix
drwxr-xr-x 14 bin        224 Nov 26 18:13 usr

# chdir /usr/source/s1

# cat echo.c
main(argc, argv)
int argc;
```

```
char *argv[];
{
        int i;

        argc--;
        for(i=1; i<=argc; i++)
                printf("%s%c", argv[i], i==argc? '\n': ' ');
}

# cc echo.c

# mv a.out newecho

# ./newecho Hello World
Hello World

# chdir /tmp

# cat >hello.c
main()
{
   printf ("Hello World!\n");
}

# cc hello.c

# ./a.out
Hello World!

# cat >hello.b
10 print "Hello World!"

# bas hello.b
run
Hello World!
```

That's it; you're up and running. You have officially set your fingers on a "real" historical UNIX system. As you can see, there is plenty of source code to look over and a working compiler to play with. UNIX V5 is only one of the early operating systems you can explore with SIMH. On the SIMH Web site, you will find a repository of disk images for other systems.

If you are interested in seeing what a PDP-11 system and RK05 disk actually looked like, take a look at the photo gallery on the SIMH Web site (see Resources). Also, try searching Google Images for a wealth of great photographs.

Resources for this article: www.linuxjournal.com/article/8587.

Matthew Hoskins is a Senior UNIX System Administrator for The New Jersey Institute of Technology where he maintains many of the corporate administrative systems. He enjoys trying to get wildly different systems and software working together, usually with a thin layer of Perl (locally known as "MattGlue"). When not hacking systems, he can often be found hacking in the kitchen. Matt can be reached at matt@njit.edu.

LINUX JOURNAL

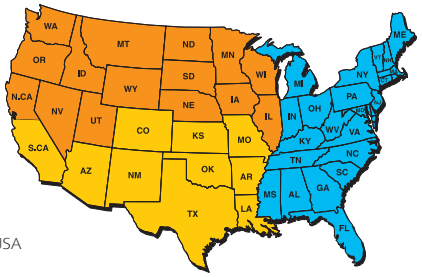

PO Box 55549
Seattle, WA 98155-0549 USA
www.linuxjournal.com

ADVERTISING SERVICES

VP OF SALES AND MARKETING
Carlie Fairchild, carlie@ssc.com
+1 206-782-7733 x110,
+1 206-782-7191 FAX

FOR GENERAL AD INQUIRIES
e-mail ads@ssc.com
or see www.linuxjournal.com/advertising

Please direct international advertising inquiries to VP of Sales and Marketing, Carlie Fairchild.

REGIONAL ADVERTISING SALES

NORTHERN USA
Joseph Krack, joseph@ssc.com
866-423-7722 (toll-free),
866-423-7722 FAX

SOUTHERN USA
Laura Whiteman, laura@ssc.com
206-782-7733 x 119

EASTERN USA
Martin Seto, mseto@ssc.com
+1 905-947-8846,
+1 905-947-8849 FAX

INTERNATIONAL
Annie Tiemann, annie@ssc.com
866-965-6646 (toll-free)

Advertiser	Page #
AML www.amltd.com	86
ApacheCon 2005 www.apachecon.com	2
APPRO HPC Solutions appro.com	55
ASA Computers www.asacomputers.com	53, 65
Cari.net www.complexdrive.com	77
Concurrent Computer Corporation www.ccur.com	17
Coraid, Inc. www.coraid.com	27
Coyote Point www.coyotepoint.com	41
Cyclades Corporation www.cyclades.com	1
D.E. Shaw Research and Development www.deshaw.com	71
Emac, Inc. www.emacinc.com	80
EmperorLinux www.emperorlinux.com	15
Etnus www.etnus.com	45
FairCom Corporation www.faircom.com	29
Free Software Foundation www.gnupress.org	63
Google www.google.com/lj	83
Hurricane Electric www.he.net	73
Iron Systems www.ironsystems.com	84
JTL Networks www.jtlnet.com	70
Kyliptix www.kyliptix.com	72
Layer 42 Networks www.layer42.net	74
Levanta www.levanta.com	7
Linux Journal www.linuxjournal.com	79, 91, 93
LPI www.lpi.org	57
MBX www.mbx.com	9
Microway, Inc. www.microway.com	C4
midPhase Services, Inc. www.midphase.com	85
Mikro Tik www.routerboard.com	13
Monarch Computers www.monarchcomputer.com	18, 19
O'Reilly Emerging Telephony Conference conferences.oreillynet.com/etel	59
Open Source Professional Institute www.ospinstitute.com	54
Open Source Storage www.opensourcestorage.com	37
Penguin Computing www.penguincomputing.com	39
PFU Systems www.pfusystems.com	87
Polywell Computers, Inc. www.polywell.com	47
The Portland Group www.pgroup.com	34, 35
Programmer's Paradise Inc www.programmersparadise.com	21, 23
Rackspace Managed Hosting www.rackspace.com	C3
R Cubed Technologies www.rcubedtech.com	49
SBE, Inc. www.sbei.com	11
SugarCRM, Inc. www.sugarcrm.com	51
Technologic Systems www.embeddedx86.com	43
Telephonyware www.telephonyware.com	25
Terra Soft Solutions, Inc. www.terrasoftsolutions.com	75
Thinkmate www.thinkmate.com	C2
TUX Magazine www.tuxmagazine.com	62
Tyan Computer USA www.tyan.com	5
USENIX lisa05	67
ZT Group International www.ztgroup.com	33

Wireless Portals with Wifidog

An easy Web-based captive Wi-Fi portal is great for users. A Web-based captive portal system that fits on a Linksys box is great for administrators too.

BY MICHAEL LENCZNER

It has become commonplace for most major cities to have a Wi-Fi group. The Wireless Community movement has spread across North America, Europe and has extended to Latin America and Asia. Hackers world-wide haven't been able to keep their hands off low-cost, easily extensible hardware. Some Wi-Fi groups get together and share technical information and war-driving data, and other groups work on projects setting up ad hoc mesh networks or creating free hotspots in their favorite hangouts.

Two years ago, in an event similar to what has taken place in many other cities, a group of Montréal technology enthusiasts got together and decided to start creating free hotspots for themselves and for other Montréalers. People joined the group after hearing about it through the local open-source grapevine. Calling themselves Ile Sans Fil (French for "Wireless Island"—and, yes, Montréal is an island), they are now one of the more active established Wi-Fi groups in the world, with 25–35 active volunteers, 50 hotspots and 6,000 users. Their current rate of expansion is 4–8 hotspots and 1,000 users per month. Based on the number of users, this volunteer group is the most successful of the seven Wi-Fi companies operating in that area.

Ile Sans Fil (ISF, www.ilesansfil.org) was able to get a quick start on the project by using a popular open-source captive portal called NoCat, which did a good job of allowing only users from a list of user names and passwords through. A captive portal is a dynamic firewall in which all traffic is blocked until the user logs in (or a disclaimer page was displayed and terms of service were agreed to). The login page works by intercepting http traffic and, in its place, displaying a form until the user is validated. Once logged in, some, or all, ports work normally. By nature, all captive portal authentication solutions are vulnerable to MAC address spoofing, and as such, these are not bulletproof. However, they have the huge advantage of not requiring any software beyond a Web browser to perform sign-on.

But NoCat wasn't perfect for their needs. The NoCat gateway was a Perl script that relied on several heavy packages. It was too big to run on most embedded hardware, so the choice was either to run it on new machines (possibly the small but expensive Soekris board) or to use old desktops dug out of closets and storage areas. Although inexpensive, the result was an open wireless access point connected to a Pentium I connected to a modem and a WAP (wireless access point). Keeping a network of heterogeneous secondhand Pentium Is running in public places proved to be a support nightmare, even for the initial three or four hotspots. The NoCat central server also lacked any network monitoring features, it was difficult to get any useful statistics from its logs and it didn't feature any mechanisms to serve different content for each hotspot. Finally, to keep a user's connection alive, NoCat used a second browser window that used JavaScript to ping the gateway every five minutes. This meant that devices that couldn't open more than one window (such as PDAs) or that had no (or disabled) JavaScript support were forced to re-authenticate continuously.

Fortunately, two years ago a wireless router running Linux became available (the Linksys WRT54G). It wasn't advertised as running Linux, but the Seattle Wireless group discovered this, and the hacking began. ISF finally had an inexpensive embedded platform to move to. They chose OpenWRT as a distro, but NoCat and its dependencies just wouldn't fit.

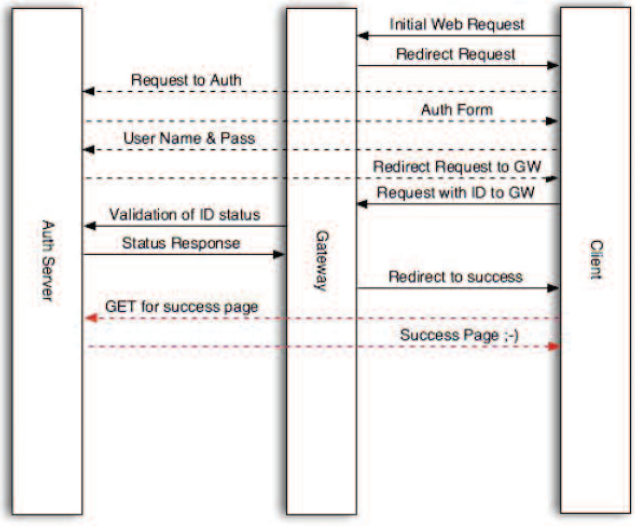

Figure 1. Wifidog has two parts: a central authentication server and a gateway located at each wireless hotspot.

And so the Wi-Fi Guard Dog captive portal system was born. Like NoCat, Wifidog (www.wifidog.org) consists of two parts: a gateway per hotspot running a client process and a Web-based central server (Figure 1).

The Wifidog gateway is written in C with no dependencies beyond the kernel. A working gateway install can be packaged in less than 15Kb on an i386 platform. It works well on the MIPS-based WRT54G running OpenWrt. The Wifidog central server is written in PHP and handles authentication and manages the network. The system knows which hotspots are up through heartbeating from the gateway.

The gateway remains small and simple by delegating all cryptography to the user's Web browser and the auth server. Tokens are generated by the auth server upon successful

Figure 2. Developers have been happy to see Wifidog adopted worldwide. From left to right: Philippe April, Pascal Leclerc, Alexandre Carmel-Veilleux, François Proulx and Benoit Grégoire. Not shown: Mina Naguib.

authentication and are then sent to the gateway. The gateway then validates the token with the auth server. Tokens are revalidated periodically in case they expire.

How secure is it? The gateway never sees the password. The token itself is transmitted in the clear between the gateway and the auth server. It would be quite simple to encrypt this, but it has been deemed unnecessary bloat, considering that it's a one-time-use token and that to do a man-in-the-middle attack on it, an attacker needs to be between the gateway and the auth server, in which case the attacker already has Internet access, making the whole attack pointless. A much more realistic attack is MAC address spoofing, which is inherently easy to do with any captive portal software running on an open Wi-Fi network. The only solution for this is to use WPA. Unfortunately, tech-support realities make it completely unrealistic to require this until every platform has a central place to enter the necessary information (not to mention that many drivers still don't support it). The team will eventually move toward 802.11i once support for the standard improves.

Of course, the Wifidog auth server handles user authentication (currently, plugins exist for internal authentication and for authenticating to a remote radius database, including logging the amount of traffic transfered by each client). But the auth server does much more than that. It handles user sign-up, real-time network monitoring, extensive statistics about network usage patterns and hotspot popularity.

With Wifidog, the volunteer group had an easy way to continue deploying hotspots while minimizing the time spent on support.

However, although this technically has been a successful project in creating another open-source captive portal solution, it is only half the story. From the beginning, ISF viewed setting up free hotspots as only a first step. The volunteers now had the tools to draw laptop users from their basements and home offices into public spaces. The next step of the project was to use the network of hotspots to help create a sense of local community.

One way in which that is done is through the promotion of local content. A unique feature of the Wifidog system is its extensive support for location-specific content. Users connecting from Café Laika see an entirely different splash page and portal page than users connecting from Atwater Library. At first, the only form that local content took was HTML and RSS feeds tied to a hotspot. Fortunately, some of the hotspots had their own RSS feeds from their Web sites.

Through working with a local new media arts group, the local content feature recently was extended, so that now there is a system that also can manage text, images, audio, video and photos from Flickr (by using the Flickr API). All of this content can be sent across the network or sent only to select hotspot portals. The extensive logging functions also allow the group to show content to a user only once, only once per hotspot, once per day. It has certainly allowed these artists some interesting and unique possibilities for location-based art.

We've got problems with your name on them.

At Google, we process the world's information and make it accessible to the world's population. As you might imagine, this task poses considerable challenges. Maybe you can help.

We're looking for experienced software engineers with superb design and implementation skills and expertise in the following areas:

- high-performance distributed systems
- operating systems
- data mining
- information retrieval
- machine learning
- and/or related areas

If you have a proven track record based on cutting-edge research and/or large-scale systems development in these areas, we have brain-bursting projects with your name on them in Mountain View, Santa Monica, New York, Bangalore, Hyderabad, Zurich and Tokyo.

Ready for the challenge of a lifetime? Visit us at http://www.google.com/lj for information. EOE

Google™

A unique feature of the Wifidog system is its extensive support for location-specific content.

Another feature is the ability to see who else is on-line at a hotspot (either locally or remotely) and find out more about them if they have filled out their profile. Profiles are an opt-in feature and not only because the group doesn't want to annoy its users. The geographical proximity of users (in the same hotspot) raises certain safety and privacy issues that don't exist in most instances of social-networking software.

This past summer has been gratifying for the developers as their project has drawn the eyes of many wireless groups all over the world. Among the groups adopting it are WirelessLondon, New York City Wireless and Paris Sans Fil. WirelessLondon has recently started to use the Wifidog gateway with their existing central server. Jo Walsh—member of the group and co-author of the recent O'Reilly book *Mapping Hacks*—writes, "We found it easy to customize for our needs; we adapted our portal service to it in half an hour. The presence of an active and committed development community around Wifidog is reassuring; we know it won't go away, and the community's been gracefully receptive to our suggestions."

Dana Spiegel—the executive director of NYCwireless—talks about his organization's impending trial of the captive portal, "NYCwireless is using the software in a pilot project and hopes to deploy it by the end of the summer to help local hotspots showcase local talent, multimedia sharing, art and student works. [Wifidog] is a great collaborative effort to provide a useful solution for community wireless networks. It enables the creation of a supported wireless network with community-oriented and created content, and really demonstrates how these networks and groups provide an important service to local areas."

The group has not been surprised by the success. Benoit Grégoire, one of the lead developers of the group, says, "We designed Wifidog to be the Swiss Army knife of captive portal systems. We hoped that it could meet the needs of most wireless community groups well enough that they would prefer to help with its development rather than roll their own. Now we're seeing some of the realization of that goal." The world of Wi-Fi community groups is starting to agree with them. What remains a question is how these other groups will use Wifidog for their own networks and in their own communities. From finding ways to make the software work (and make sense) in a mesh network, to developing GIS applications, to adding chat functionality to the network, there's lots of promising community and social applications for what was originally an infrastructure project.

Beyond the interesting technical possibilities, it is the chance to have an impact on the lives of their fellow citizens that seems to motivate Wifidog developers the most. With 10,000 users expected by December 2005 in Montréal alone, there is a good chance that their code will be used by neighbors, coworkers and friends. That, combined with the frequent press coverage and the chance to work with people they wouldn't normally meet, such as artists and community activists, means the team's energy and enthusiasm should remain high for the foreseeable future.

Michael Lenczner is a volunteer with Ile Sans Fil. He has been working in community informatics for eight years, both in Canada and abroad. He blogs at mtl3p.ilesansfil.org.

Ultra Dense, Powerful, Reliable...
Datacenter Management Simplified!
15" Deep, 2-Xeon/Opteron or P4 (w/RAID) options

Customized Solutions for...
Linux, BSD, W2K
High Performance Networking Solutions
• Data Center Management
• Application Clustering
• Network and Storage Engines
Rackmount Server Products
• 1U Starting at $499: C3-1GHz, LAN, 256MB, 20GB IDE
• 2U with 16 Blades, Fast Deployment & more...

iron SYSTEMS™
Iron Systems, Inc.
2330 Kruse Drive, San Jose, CA
www.ironsystems.com
CALL: 1-800-921-IRON

Vim for C Programmers

You don't have to move to an integrated development environment to get luxury coding features. From variable autocompletions all the way up to integration with ctags and make, Vim makes a C programmer's life easier and more productive. BY GIRISH VENKATACHALAM

Vim is an extremely powerful editor with a user interface based on Bill Joy's almost 30-year-old vi, but with many new features. The features that make Vim so versatile also sometimes make it intimidating for beginners. This article attempts to level the learning curve with a specific focus on C programming.

make and the Compile-Test-Edit Cycle

A typical programmer's routine involves compiling and editing programs until the testing proves that the program correctly does the job it is supposed to do. Any mechanism that reduces the rigor of this cycle obviously makes any programmer's life easier. Vim does exactly that by integrating **make** with Vim in such a way that you don't have to leave the editor to compile and test the program. Running `:make` from inside of Vim does the job for you, provided a makefile is in the current directory.

You can change the directory from inside of Vim by running `:cd`. To verify where you are, use `:pwd`. In case you are using FreeBSD and want to invoke gmake instead of make from the command line, all you have to do is enter `:set makeprg=gmake`. Now say you want to give some parameters to make. If, for instance, you want to give CC=gcc296:

```
:set makeprg=gmake\ \CC=gcc296
```

does the job.

Now comes the job of inspecting the errors, jumping to the appropriate line number in the source file and fixing them. If you want to display the line numbers in the source file, `:se nu` turns on this option, and `:se nonu` disables line number display.

Once you compile, Vim automatically takes you to the first line that is causing the error. To go to the next error; use `:cn` to take you to the next line number causing the error. `:cfirst` and `:clast` take you to the first error and the last error, respectively. Once you have fixed the errors, you can compile again. If you want to inspect the error list again, `:clist` displays it. Convenient, isn't it?

If you want to read some other source file, say foo.c, while fixing a particular error, simply type `:e foo.c`.

One shortcut provided by Vim to avoid typing too much to switch back to the previous file is to type `:e #` instead of typing the full path of the file. If you want to see all of the files you have opened in Vim at any point in time, you can use `:ls` or `:buffers`.

If you have a situation in which you have opened too many files and you want to close some of them, you can issue `:ls`. It should display something like this:

```
2 #     "newcachain.c"                    line 5
3 %a    "cachain.c"                       line 1
```

If you want to close newcachain.c, `:bd 2` or `:bd newcachain.c` does the job.

While browsing C code, you may have situations in which you want to skip multiple functions fast. You can use the]] key combination for that while in command mode. If you want to browse backward in the file, [[can be used.

You also can use marks to bookmark certain cursor positions. You can use any lowercase alphabet character as a mark. For instance, say you want to mark line number 256 of the source and call it b. Simply go to that line, `:256`, and type `mb` in command mode. Vim never echoes what you type in command mode but silently executes the commands for you.

If you want to go to the previous position, typing `''` (two single-quotation marks) takes you there. Typing `'a` takes you to mark a and so on.

Especially when editing Makefiles, you may want to figure out which of the white spaces are tabs. You can type `:se list`, and whatever is displayed as `^I` in blue are tabs. Another way to do that is to use `/\t`. This highlights the tabs in yellow.

Global searches and replaces are common tasks for programmers, and Vim provides good support for both. Simply type `/` in command mode, and you are taken to the searched keyword. If you prefer incremental searches, à la emacs, you can specify `:se incsearch` before you search. When you want to disable it, type `:se nois`.

Search and replace is a powerful tool in Vim. You can execute it only on a region that you selected using the v command, only between certain line numbers or only in rectangular regions selected by using Ctrl-V command.

Once you select your region or line number ranges, for example using `:24,56` to select lines 24–56 (both inclusive), type `s/foo/bar` to replace all occurrences of the string foo with bar.

But, this command replaces only one instance per line. If you want to do this for multiple occurrences per line, type `s/foo/bar/g`. If you want to replace only some occurrences, you can use the "confirm" option with `s/foo/bag/gc`.

Sometimes the string contains characters that appear as a substring of other keywords. For instance, say you want to replace the variable "in" and not the "in" in inta. To search for whole words, type `/\<in\>/`.

Most commonly, you will want to do a global replace, which is every instance in a given file. You can do that by using either `:1,$s/foo/bar/g` or `:%s/foo/bar/g`. If you then want to replace this in all the files you have open, you can enter `:bufdo %s/foo/bar/g`.

Another way of searching is by going to the keyword and typing * in command mode. The keyword now will be highlighted wherever it occurs in the file. Searching backward is simple too; type ? instead of / while searching.

Once the searching is over, Vim remembers it, so the next time you search for the same keyword, you have to type only / or ?, instead of typing the whole text.

One side effect of searching is that it stays highlighted. This can be a distraction while editing programs. Turn highlighting off by typing `:se nohlsearch`, `:nohlsearch` or `:nohl`

You always can use the Tab key to complete Vim commands you give with a colon. For instance, you can type `:nohl`<Tab>, and Vim completes it for you. This is applicable generically, and you can press Tab to cycle through Vim's commands until Vim finds a unique match.

Vim with Exuberant ctags

Exuberant ctags (see the on-line Resources) is an external program that can generate tags for Vim to navigate source code. If all of your source code is contained in only one directory, simply go to the directory in the shell and enter:

```
$ ctags .
```

This generates a tags file called tags. Vim reads this file for jumping to functions, enums, #defines and other C constructs.

If the source code is distributed across several directories, ctags has to generate tags for all of them relative to a certain directory. To do this, go to the root directory of the source code and execute:

```
$ ctags -R .
```

Check whether the tags file has been generated. You also can open and read the tags file in Vim.

Now, let us move on to navigating the source code using tags. Navigating the source code using ctags is one of the most fascinating tools that a programmer has. You can read the code so nicely and quickly that you wonder how it would have been without ctags.

Once the tags file has been generated, open the file in Vim as normal, except that if the file is deep inside, open it

from the root directory. For instance, your source code is organized like this:

```
common
     |
     ----> gui --> wxpython
     |        |
     |        ------>Tk
     |
     ----> backend --> networking

  include

  user
```

If you want to edit tcp.c under the common/backend/networking directory, you should open it like this:

```
$ vim common/backend/networking/tcp.c
```

instead of like this:

```
$ cd common/backend/networking
```

and:

```
$ vim tcp.c
```

The tags file is situated in the directory above common, and Vim automatically knows the location of the tags file this way.

Alternatively, you can open the file using the second method mentioned above and execute this from inside of Vim:

```
:se tags=../../../tags
```

The first method is easier for navigation. Once you open the file, you can jump from one function definition to another easily by using the key combination Ctrl-].

If you want to go to the definition of anything, be it a function, macro or anything else, simply press Ctrl-] when the cursor is positioned on it. Thus, from invocation, we can move to the definition. It takes you there no matter which file contains it. Assuming that we call drawscreen() from tcp.c, it automatically takes you there, even if the file is contained under common/gui.

If you want to go back to what you were reading, press Ctrl-T, and you return to where you left. You can jump to another invocation from there by pressing Ctrl-] again. You can continue this process ad infinitum, and you can keep coming back by pressing Ctrl-T.

Another way to find a function definition if you know only a part of the name is:

```
:ta /function
```

This command takes you to the first match if there are multiple matches. You can go to the next match with `:tn`.

If there are multiple definitions and you want to choose among them, you can press G Ctrl-] or type `:tselect <tagname>`. This way you can modify the source code by navigating with tags without even knowing which file contains what.

Vim with cscope

cscope is another powerful source code navigation tool with which we can perform a variety of searches. Here is a sample output of the cscope menu:

```
Find this C symbol:
Find this global definition:
Find functions called by this function:
Find functions calling this function:
Find this text string:
Change this text string:
Find this egrep pattern:
Find this file:
Find files #including this file:
```

Now, Vim has integrated cscope into its repertoire, making it convenient for programmers to use the same features in cscope from the cool comfort of Vim. All you have to do is establish a cscope connection by issuing `:cs add cscope.out`.

As we discussed before with ctags, cscope generates an index called cscope.out that can be generated by using the shell command:

```
$ cscope -Rbq
```

This generates the file cscope.out. It is to be executed from the source code root directory à la ctags. You then open the file as

before, relative to the source code root directory, and make a cscope connection with the command `:cs add cscope.out`. You can verify existing cscope connections by typing `:cs show`.

What you can search for from inside of Vim can be seen using `:cs<CR>`. For instance, to go to a particular file, or a header of a source file, simply type `:cs f f stdio.h` for opening stdio.h or `:cs f f foo.c`.

For searching for functions called by a function foo.c, type `:cs f d foo.c`. This lists out the functions called by foo.c. For functions calling foo.c, type `:cs f c foo.c`.

To search for an egrep pattern, type `:cs f e varName` and so on. For a list of the available options, type `:cs`. It displays a range of available options.

Now, if you have both ctags and cscope, you can type `:cstag /foo` to search for a function or enum or whatever that contains foo.

Vim and Syntax Highlighting

If there is one feature in Vim for which it wins hands-down compared to any other editor or IDE, it is full-featured syntax highlighting. The colors available in Vim make it a veritable delight to work with source code. It not only makes your life colorful, it also makes it easy to spot errors ahead of compilation. Common errors such as a mismatched),} or] in the code are easy to see. It also reminds you if you have left a string hanging without the closing " or '. It tells you the comment doesn't end with */, or that you are nesting comments. Syntax highlighting is smart when it comes to C syntax.

Typically, you wouldn't have to do anything to enable Vim's syntax highlighting; `:sy on` does the job in case your distribution doesn't enable it by default. As with other commands, you can add this to your ~/.vimrc file.

If colors still don't show up, something is wrong with your terminal. Fix it first. `:se filetype on` is another thing you can try in addition to `:sy enable`.

Let us assume that you have colors displayed correctly. Say you don't like a certain color, or because blue is not visible in dark backgrounds, you can't read C comments. To solve the second problem, a simple `:se background=dark` does the job. If you want to disable syntax highlighting for C comments, type `:highlight clear comment`.

To change colors, first use the `:syntax` command to display all the syntax items for the given buffer. Then, identify the syntax group you want to change. If you want strings displayed in a bright white color, which is easy to read against a black background, simply enter:

```
:highlight String ctermfg=white
```

or, for gvim users, type:

```
:highlight String guifg=white
```

You also can change the syntax color of any group. Typical syntax groups are Statement, Label, Conditional, Repeat, Todo and Special. You can change the attributes of highlighting as well, such as underline and bold. For instance, if you want to display NOTE, FIXME, TODO and XXX with underlining, you can use:

```
:highlight Todo cterm=underline
```

or you can both add bold and change the color:

```
:highlight Repeat ctermfg=yellow cterm=bold
```

You can create your own set of highlight commands and save it in your ~/.vimrc file so that every time you edit your source code, your favorite colors are displayed.

Vim and Variable Name Completion

In addition, Vim has a feature for variable name completion. While typing, simply press Ctrl-N or Ctrl-P in insert mode. Remember, this works only in insert mode. All other commands mentioned above work in command mode. You can cycle through possible completions by pressing Ctrl-N again.

This helps us avoid errors while typing, because structure members and function names often can be misspelled. This works best when Vim can use tags, so make sure a ctags file is in place.

Vim and Source Code Formatting

Vim understands C well enough to be able to indent code automatically. The default indentation style uses tabs, which may not be appropriate for some people. In order to remove tabs completely from the source, enter:

```
:set expandtab
:retab
```

which converts all tabs into spaces in such a way that the indentation is preserved. While typing C text, Vim automatically indents for you. This helps you figure out where you have your matching brace. You can match braces,),] and } with the % command in command mode. Simply take the cursor to a brace and press %, which takes you to the corresponding closing or opening brace. This works for comments as well as for #if, #ifdef and #endif.

After finishing typing the program, if you want to indent the whole file at one go, type `gg=G` in command mode. You then can remove tabs if you want by the above-mentioned method. `gq` is the command sequence for indenting comments. You can select a region and indent it too with the = operator.

If Vim's default tab indentation is painful to use, you can disable it by setting `:se nocindent`. Other indentation options are available. You can indent code between two braces and between certain line numbers. You can learn more by typing `:help indent.txt`.

Conclusion

Vim comes with rich help documentation. Type `:help` from inside of Vim to browse it. To go to a particular topic, press Ctrl-] on the turquoise-colored text. Vim's help documentation uses the navigation mechanism we saw using tags.

Resources for this article: www.linuxjournal.com/article/8455.

Girish Venkatachalam loves to play with open-source operating systems, such as OpenBSD, FreeBSD and Debian GNU/Linux. He also likes to go cycling when not hacking. He can be contacted at girish1729@gmail.com.

Mini KDE for a Lightweight Desktop

Do you need a memory hog of a desktop environment simply to run a few essential programs? This experiment says you might not.

BY MARCO FIORETTI

Many users need computers only for basic office productivity, Web access and e-mail. Free software for all of these tasks exists, but it has a hidden cost. Often, students, schools and charities can afford only hardware that is five or more years old, with limited CPU power and disk space. As weird as it seems, the latter often is the most serious, apparently unsolvable problem. You may need only five or six small programs, but they are available only in big bundles, which in turn have many more dependencies. The real, total space requirements can be heavy enough to make the installer abort for lack of space.

Often, installing current but feature-light applications is useless. Desktop computers are communication tools. Today, that means, at least, digital signatures, IMAP support, checking one's bank account by way of SSL or XHTML Web forms and so on. It also means support must be provided for OpenDocument, an office file format, default in OpenOffice.org 2.0, that has raised great interest in the European Union and soon will become an ISO standard.

Installing older distributions is useless for the same reasons and is dangerous to boot: why would people go on-line and expose themselves to a bunch of security holes that have been known about for years? Furthermore, free on-line support for five- and six-year-old code is practically nonexistent, unless you have the time and skills necessary to hack together a fix for yourself.

All this is why, a few years ago, I and others started the RULE Project—to make it possible to use old hardware with current, mainstream GNU/Linux applications by installing only what truly is needed. Our approach, however, offers several advantages to modern hardware as well. First, the RULE Project makes it easier to run any computer at its greatest possible speed.

The second advantage is running normal x86 software with something built today that is much smaller and less power-hungry than a laptop. Last year, a user working to make a desktop box out of a Norhtec Microclient wrote that he "was delighted to see that RULE provides ALSA, Udev and all the other up-to-date goodies...in only 232MB...because Fedora 3 provides them".

The third big stimulus to trim down modern programs also has nothing to do with vintage computers: bootable Linux CD-ROMs and USB drives are great as portable emergency desktops but require little space.

There is one final reason why all this exercise is worthwhile, but it is of interest only to KDE developers and packagers, so I'll mention it later.

Project Specifications

What are the characteristics of a useful yet lightweight desktop? To me, they are the ones just mentioned. This is why I decided to re-package together KOffice, Konqueror, KMail, KNode and almost nothing else.

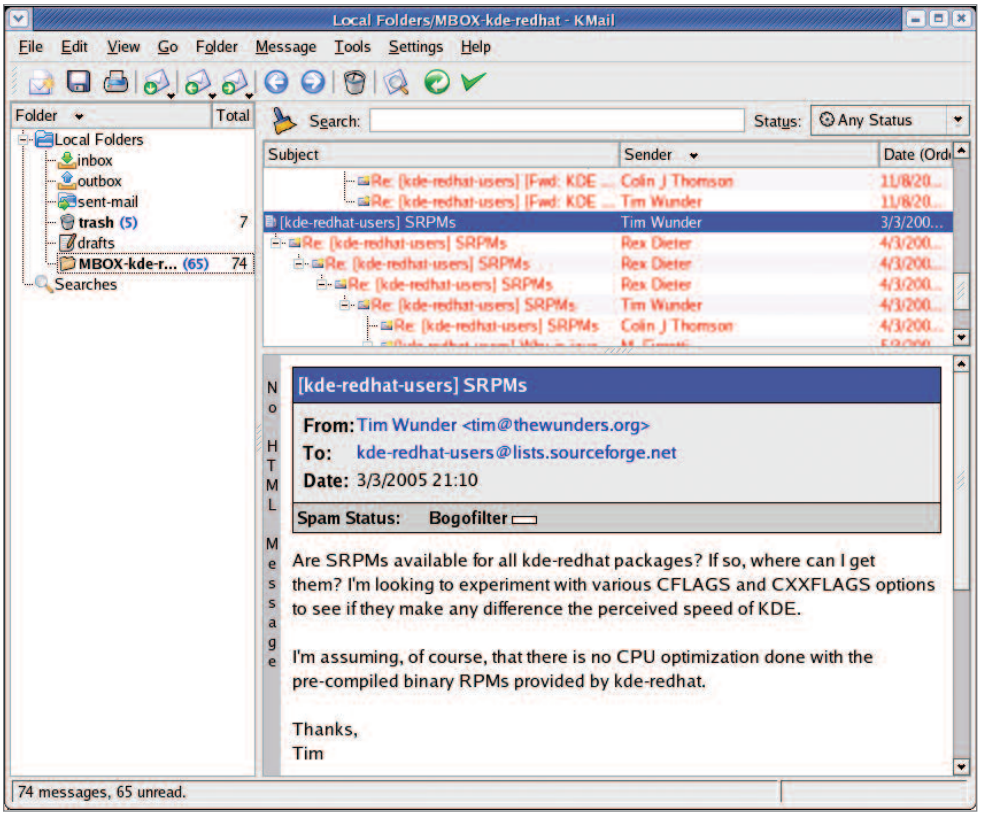

Figure 1. KMail for Mini KDE

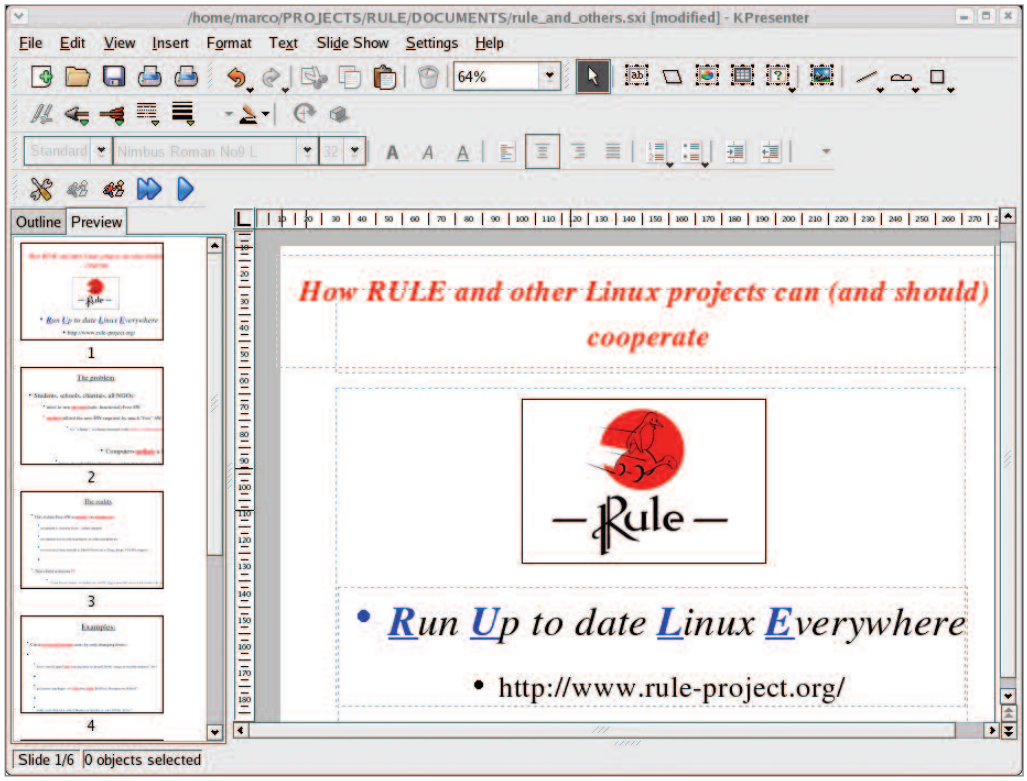

Figure 2. KOffice for Mini KDE

KOffice does not have as many features as does OpenOffice.org, but it is much lighter, is less reliant on Java, is more integrated with Linux and could, some day, share single-file SQL databases with OpenOffice.org (see the on-line Resources). Above all, KOffice's roadmap officially foresees full support for OpenDocument. The result, which we hereby call Mini KDE, must require the smallest possible disk space and RAM to run. The rest of this article summarizes what I did to achieve this goal.

How Can We Do It?

I wanted to end up with binary packages, because many desktop end users don't know how to compile by themselves, and it would be time consuming to do it on six- or seven-year-old boxes (if not impossible, because compiler, libraries, source code and intermediate compilation files would, again, *not* fit on a smaller hard disk). Generally speaking, one can obtain optimized KDE packages by using three different methods:

1. Optimize the source code of the application(s) and related libraries with the proper compiler options.

2. Compile, package and install only selected pieces of the bundle.

3. Configure the result so that applications start and run more quickly.

The last method can or must be applied even after installation. For KDE applications, it already is documented in the KDE performances tips page (see Resources).

The first method is distribution- and compiler-dependent; it's also beyond the skills of nonprogrammers such as myself and most general users. Another problem is almost no related information is available on-line; even asking on developers lists didn't turn up much more help. Carried to the extreme, this method also implies compiling against a custom version of Qt, stripped as discussed on the RULE Web site, which is almost like creating yet another distribution. From my point of view, however, the biggest limit of this method is that it does not greatly reduce the size of the whole package, which we saw as the first obstacle.

The most promising strategy, and the one I discuss in the rest of this article, is the second one—to leave out as much as possible from the original bundles in a way that minimizes effort, required skills and risk. The explanations that follow are based on building RPMs for Fedora 3, but the general procedure is valid for every GNU/Linux distribution or packaging format. Apart from the biggest space savings, another great advantage of this method is the resulting binaries remain compatible with Fedora Core or whichever other mainstream distribution you started with.

Preparation

First of all, I cleaned up my computer running Fedora Core 3. Partly, this was done to make some extra room, but the main reason was to build the packages in a clean environment. After some checking and thinking, I removed the following packages, which I originally had installed from Fedora Core or KDE/Red Hat repositories: kdeedu, kdeartwork, KOffice, kdesdk, kdevelop, kdepim, kde, kdebase, kdelibs and kdelibs-devel.

Here's the other reason to perform such trimming exercises: you can learn a lot about how packages relate to one another. Specifically, you discover unneeded dependencies and packaging errors that remain hidden when distributions simply bundle software together without paying attention. For example, I learned that, at least on Fedora, I couldn't remove redhat-menus-3.7.1-3.4.3.kde, because it is needed by apparently unrelated stuff, including htmlview, gnome-vfs, openoffice.org-1.1.2, Evolution, XMMS and Nautilus.

The same happened with arts, the modularized sound system for KDE, and its development complement, arts-devel. Users of older desktops certainly are able to survive, even

when they have a sound card, without acoustic effects. However, those two packages are needed by many more applications, including gstreamer plugins, gnome-applets, Evolution and so on. Some of these dependencies do make sense once you find them, but others still make me wonder. In any case, there seems to be a lot of opportunities for space savings at this level.

After cleaning my hard disk, I installed the latest stable source RPMs of kdelibs, kdebase, kdepim and KOffice from apt.kde-redhat.org/apt/kde-redhat/all/SRPMS.stable. When I started, they were:

- kdebase-3.4.1-1.0.kde.src.rpm
- kdelibs-3.4.1-1.0.kde.src.rpm
- kdepim-3.4.1-1.1.kde.src.rpm
- koffice-1.3.5-3.0.kde.src.rpm

I chose the KDE for Red Hat Project instead of official Fedora Core packages, because I find them more polished than the standard ones. They also usually offer newer versions of the packages.

How I Did It

When you install a source RPM, you get all of the source code in a .tar.bz2 archive and the instructions to build everything in a .spec file. Normally, to build the package, you need to issue only the command:

```
rpmbuild -ba <package_name>.spec
```

To reduce disk space, I basically did two things, both relatively simple even for nonprogrammers. The first was to massage the compile and installation options in the .spec files. For example, I compiled everything without sound, adding `-without-arts` to the configure section. When available, I also added similar options to ignore other multimedia libraries or support for devices such as cell phones and PDAs. Then, I commented out all the Require and BuildRequires directives that check whether libraries for audio, video and modern peripherals are available before starting the process. I also removed the Provides directives for all the binaries I left out. Finally, I commented out the instructions that pack into the binary RPM files that I had not compiled or didn't need.

My complete .spec files are available in the Mini KDE section of the RULE Web site.

The second and most important trick was to insert a proper inst-apps file inside each KDE source tarball. It turns out that the configure scripts of these programs have a section that more or less says something like this (from kdelibs):

```
ac_topsubdirs=
if test -s $srcdir/inst-apps; then
  ac_topsubdirs="`cat $srcdir/inst-apps`"
elif test -s $srcdir/subdirs; then
  ac_topsubdirs="`cat $srcdir/subdirs`"
fi
```

$ac_topsubdirs is the list of all the subdirectories whose code must be compiled and installed. By default, this variable is loaded with everything written in the subdirs file. But, if you copy subdirs into inst-apps, remove from the latter all the unneeded items and then tar and compress everything again, only the applications you want are compiled. This also works when installing directly from source.

Generally speaking, to figure out what you could or could not remove from inst-apps, look at the README file in each subdirectory. The following is a short summary of what I did for each package.

kdelibs

I removed only the following items: arts, kdoctools, kate, libkscreensaver and doc. In the %configure section, I excluded

Figure 3. Konqueror for Mini KDE

xinerama, alsa and artsd support. I also commented out the Requires: arts directive, as well as those for jasper and openexr.

kdebase

The only pieces I wanted from kdebase were libkonq, Konqueror, Kicker and Kwin. I was able to exclude support for xinerama, jasper, arts, Java, GL, Samba, lm-sensors, mDNSResponder-devel and libraw1394-devel. I left out the wallpapers. I also removed sounds and templates, together with the dependency from the redhat-artwork package. But, I had to put them back, otherwise RPM couldn't make it to the end for reasons not clear yet.

kdepim

Here, as I needed only KMail, Kopete and KNode, I removed a lot of programs: karm, knotes, kdgantt, kgantt, korn (mail notifier), kpilot, kmobile and ksync, kandy, kitchensync, kalarm, kresources, kfile-plugins, konsolekalendar, korganizer, wizards, kontact and plugins. Even the BuildRequires dependencies from bluez-libs-devel (Bluetooth) and gnokii (Nokia phone support) went away without problems.

KOffice

Nothing was done here, except for the addition of the -without-arts configure directive.

Final Results

Table 1 shows the sizes of the resulting binary packages, the first column, compared with standard RPMs for the same source versions from Fedora Core 4 or, for KOffice, Fedora Core 3 update repositories.

To summarize, I went from a total of 78.24MB to 57.29MB for the four packages above. This is a 26.8% reduction in file size, which doesn't look bad at all, but the final space savings was only 20.95MB. The actual impact on disk space is better, however; Mini KDE required a bit less than 150MB. The regular packages for the same four bundles, plus the extra ones they carried along, came to just less than 340MB.

Keep in mind, these are my results from only the first trial, without changing or ever looking at the source code and maintaining full compatibility with my chosen distribution, all its updates and any third-party Qt programs. All the screenshots in this article show that the resulting binaries run without problems on Fedora Core 3.

You probably noticed that the only real savings come from kdebase and kdepim. This was expected. I haven't found out yet why KOffice came out slightly bigger, but I wanted to keep functionality, so I didn't remove anything from it. I simply

Table 1. Binary Package Size Comparison (sizes in bytes)

Package Name	Mini KDE	Fedora
kdebase-3.4.1-1.0.kde.i386.rpm	17,798,755	27,736,762
kdelibs-3.4.1-1.0.kde.i386.rpm	15,109,882	18,140,844
kdepim-3.4.1-1.1.kde.i386.rpm	9,864,436	18,089,962
koffice-1.3.5-3.0.kde.i386.rpm	14,514,826	14,276,427

Do you take

"the computer doesn't do that"

as a personal challenge?

So do we.

LINUX JOURNAL™

Since 1994: The Original Monthly Magazine of the Linux Community

Subscribe today at www.linuxjournal.com

rebuilt the package to make sure that my reduced kdebase and kdelibs were compatible with it.

As far as the other packages go, KDE is a bundle of many programs built on a common foundation. Even if you use few programs, that set of core libraries, dæmons and what-not cannot become much smaller. This is why kdelibs and part of kdebase remained almost untouched. At the same time, saying "I want only five or six applications, not 40" is what actually made kdebase and kdepim much smaller, almost without affecting the functionality of the remaining programs.

Conclusion and Credits

There are surely things that I have missed, tricks that I still have to learn and space for a lot more improvement in the method I have described. However, this was only a first test: the final goal, besides reducing the package size, is to make the compilation and packaging process of this Mini KDE as automatic as possible on every distribution. In this way, whenever new KDE or KOffice versions are released, they quickly and easily could be made available to all users with limited hardware and not enough skills to start from the source.

In order for this to happen, it is necessary to discover, collect and write down as much information as possible on how the items in the several subdirs files are related to one another, as well as any other optimization tricks. Suggestions are welcome!

I will continue to experiment in this area with the folks of the RULE and Ubuntu-lite mailing list, which I thank for their support and interest in this idea. You can find all the results and instructions for Mini KDE on the RULE Web page.

Special thanks also go to Luciano Montanaro, D. Faure and all the KOffice developers who provided much of the initial information to get me started.

Resources for this article: www.linuxjournal.com/article/8536.

Marco Fioretti is a hardware systems engineer interested in free software both as an EDA platform and, as the current leader of the RULE Project, as an efficient desktop. Marco lives with his family in Rome, Italy.

Might Be Just Right

At LinuxWorld in Boston earlier this year, I got together with an old Swedish friend. She's a nurse, not a technologist, but she was curious about my work and the conference that brought me to town. Somewhere in the midst of my explanation of Linux and its virtues, she said, "Ah, Linux is lagom." She explained that *lagom* is a Swedish term that conveys a sense of balance, proportion and appropriateness. "Not too much, not too little...just right."

When I told her that Linus Torvalds' first language and surname were both Swedish, she said, "Well of course. There you go." (I'm half-Swedish myself, though I'm not sure that matters.)

So I put the question "Is Linux logom?" to The Man Himself in an e-mail. He debugged my spelling and declined to commit:

> Lagom, with an "a".
>
> And yes, it means "just right", in the sense of "not too much, not too little". See en.wikipedia.org/wiki/Lagom

Then he added, in a following e-mail:

> They still end up confusing "lagom" with finding the "optimal" amount. That's pretty much missing the point. It's not that something is "lagom" because it's the best possible or "optimal". Quite the reverse. Something being "lagom" very much involves not caring too much about what the optimal amount even is. Or possibly questions where "optimal" simply doesn't make sense.

So I began checking other sources. The best I found was from "In Other Words", published in AskOxford, published by the Oxford English Dictionary (www.askoxford.com/worldofwords/wordfrom/otherwords). It lists lagom among a handful of "the most insightful, intriguing, and satisfying expressions on the planet—for which there are no English equivalents". It says:

> Swedish commentator Dr Bengt Gustavsson argued that the lagom mentality can be seen as the trait that gives Swedish society its characteristic stability and yet an openness to external influences. The word alludes subconsciously to the avoidance of both conspicuous success and humiliating failure, which is deeply ingrained in the Swedish psyche. It is the inclination among Swedes to shun ostentation, accept modest rewards, be good team players—to fly beneath the radar.

FREE NEWSLETTER!

Wish you could get the latest from *LJ* more than once a month? You can—sign up today for ***LJ*'s weekly e-mail newsletter**.

Each week the *LJ* newsletter features great tech tips, links to web-only articles, and news on the latest events in the Linux market.

Sign up for the *LJ* e-mail newsletter now:

http://www.linuxjournal.com/

NetOp RemoteControl

Secure Remote Control & Support for Linux

Award-winning NetOp Remote Control for Linux provides secure, cross-platform, remote control, access and support. NetOp lets you view and control a remote PC's current desktop session, transfer and synchronize files, launch applications or chat with the remote user - just as if you were seated at that computer.

> Cross-Platform support for Linux, Solaris, Mac OS X, & all Windows platforms
> Advanced security including encryption, multiple passwords, even centralized authentication & authorization with the optional NetOp Security Server module

NetOp and the red kite are registered trademarks of Danware Data A/S. Other brand and product names are trademarks of their respective holders. 2003 Copyright Danware Data A/S. All rights reserved

Try it Free - www.CrossTecCorp.com CrossTec Corporation

Who says penguins can't fly?

http://store.linuxjournal.com/

STOP SPAM

CALL TODAY
And start selling the most powerful anti-spam solution to your clients: CanIt-PRO

PARTNER WITH ROARING PENGUIN

Roaring Penguin is looking for a few good resellers for Can-It PRO:
- The most flexible anti-spam solution on the market
- The easiest to resell
- Ideal for Linux consultants

ROARING PENGUIN SOFTWARE INC.

www.roaringpenguin.com/partners (613) 231-6599

www.linuxjournal.com

Low Total Cost Of Ownership!

Fitrix Exact Fit Accounting & Business Software

Source Code Accounting & Distribution Software For Linux

Free Trial!

www.fitrix.com
800.374.6157
770.432.7623

You'll be THANKFUL for all the time and money you save with FITRIX!

Open-Source Use Accelerates Software Development

Best practices for managing software license compliance in an Open Source world. **BY PALLE PEDERSEN**

Many software developers share some common traits. Besides copious caffeine and creative work schedules, developers like working on interesting, new problems. They like to be as efficient and productive in their work as possible. They never want to start from scratch either; they prefer to cut, paste, modify and extend.

The best developers today are the ones that can most effectively find, assemble and then optimize re-usable software components—whether those components are open source or were developed within their own organization. A developer's skill with Google and SourceForge is now as important as his or her knowledge of software architecture and implementation.

This new "assembly" model is fundamentally changing the way software is designed and developed. It accelerates development, improves software quality and reduces costs. In short, it's changing everything.

Software now consists of a mix of company-owned code, open-source and commercial libraries, and code provided by outsourcers. By combining external components with their proprietary technology, companies create a set of mixed intellectual property, or mixed-IP assets. Best practices for managing software licensing in this new mixed-IP environment are now emerging.

All software—whether commercial or open source—carries licensing obligations that companies must comply with. This new, mixed-IP environment adds complexity to the process by mixing together licenses of all kinds. Managing these licenses and their restrictions needs to be done correctly in order to keep companies—and individual developers—out of trouble.

Following is a set of software compliance management "best practices" that were developed through discussions I've had with companies that are best in the world at leveraging this new environment:

1. Re-use existing components—to lower development costs, accelerate time to market, improve quality and reduce business risk, use existing internal and external components wherever appropriate. Explicitly consider functionality, performance, reliability, maturity, risk, sensitivity and license obligations.

2. Track and control changes to internal components—to establish and maintain the provenance of all internal components, to identify and protect critical IP and to avoid inadvertent violations of licenses, trademarks, patents, copyrights and trade secrets. Track internal component creation and modification and control the modification of those that are sensitive.

3. Control re-use of sensitive or external components—to avoid last-minute surprises, guesswork, compromises and risk-taking, and to prevent the loss of intellectual property and facilitate timely and effective remediation. Review and approve the use of any external or sensitive internal components or fragments in a project.

4. Verify every build and release—to assure prompt discovery of materials inadvertently included in a project and unapproved or precluded modifications to components. Identify and remediate all unapproved components or fragments and changes made to any of those components. Record the metadata for all external components in the associated bill of materials.

5. Review compliance at project phase transitions—to prevent loss of intellectual property and to assure prompt discovery of new components inadvertently included in the project. At major development milestones, verify that no unapproved components are used in the project or were changed and then used. Review the license obligations of all external components used in the project and ensure compliance with them.

6. Control component contribution and disposition—to avoid license violations and the attendant disruptions and to constrain the propagation of risky software. Before contributing any component or fragment to an open-source project or transferring ownership to another party, assess the sensitivity of that material. Verify your rights to make that contribution or transfer.

7. Assess software components before acquisition—to prevent negative post-acquisition surprises. Before buying a software component, identify all internal and external components used in that asset. Identify all external components used in any active project and assess their license obligations with respect to compliance, business objectives and legal policies. Assess the impact of any required rework or change on cost, revenue, quality and so forth.

These best practices encourage the use of open source and re-use of software components, while assuring compliance with license obligations. They also protect an organization's intellectual property assets. In addition to adopting these best practices, many organizations are using commercially available automated solutions as platforms on which to establish and manage these processes.

Open-source software and component re-use are here to stay. Now is the time for companies to begin thinking about how they will alter their management of software IP so they can capitalize on this new development model. By doing so, they will get ahead of the issue and put the power of open source to work for their organizations.

Palle Pedersen is CTO of Black Duck Software, the leading provider of software compliance management solutions (www.blackducksoftware.com).

Rackspace — Managed Hosting backed by Fanatical Support.™

Servers, data centers and bandwidth are not the key to hosting enterprise class Web sites and Web applications. At Rackspace, we believe hosting is a service, not just technology.

Fanatical Support is our philosophy, our credo. It reflects our desire to bring responsiveness and value to everything we do for our customers. You will experience Fanatical Support from the moment we answer the phone and you begin to interact with our employees.

Fanatical Support has made Rackspace the fastest-growing hosting company in the world. Call today to experience the difference with Fanatical Support at Rackspace.

LINUX JOURNAL READERS' CHOICE 2004

Thanks for honoring us with the 2004 Linux Journal Readers' Choice Award for **"Favorite Web-Hosting Service"**

rackspace™ MANAGED HOSTING

1.888.571.8976 or visit us at **www.rackspace.com**

The World's Fastest InfiniBand Switch

From a Company You've Trusted for 23 Years

▲ 72 Port FasTree™ Configuration

Microway's **FasTree**™ DDR InfiniBand switches run at 5GHz, twice as fast as the competition's SDR models. FasTree's non-blocking, flow-through architecture makes it possible to create 24 to 72 port modular fabrics which have lower latency than monolithic switches. They aggregate data modulo 24 instead of 12, improving nearest neighbor latency in fine grain problems and doubling the size of the largest three hop fat tree that can be built, from 288 to 576 ports. Larger fabrics can be created linking 576 port domains together.

Working with PathScale's InfiniPath HTX Adapters, the number of hops required to move MPI messages between nodes is reduced, improving latency. The modular design makes them useful for SDR, DDR and future QDR InfiniBand fabrics, greatly extending their useful life. Please send email to fastree@microway.com to request our white paper entitled *Low Latency Modular Switches for InfiniBand.*

Microway's **QuadPuter**® includes four AMD single or dual core Opteron™ processors, 1350 Watt redundant power supply, and up to 5 redundant, hot swap hard drives-all in 4U. One of the most powerful processing platforms in the HPC industry, QuadPuter can serve as a cluster node or a standalone supercomputer. Constructed with stainless steel, its RuggedRack™ architecture is designed to keep the processors and memory running cool and efficiently. The power supply exhaust does not mix with air in the motherboard chamber. Hard drives are cooled with external air and are front-mounted along with the power supply for easy access and removal. The RuggedRack™ is available with an 8-way motherboard, dual-core Opterons and up to 128 GB of memory for power- and memory-hungry SMP applications.

◀ QuadPuter® Navion™ with Hot Swap, Redundant Power & Hard Drives

Call us first at 508-746-7341 for quotes and benchmarking services. Find technical information, testimonials, and newsletter at microway.com.

Visit us at SC2005 Seattle

AMD 64 Opteron™

PathScale™ Accelerating Cluster Performance™

Microway®
Technology you can count on℠

508.746.7341 microway.com

8-Way Opteron Server

Introducing the Thinkmate 5U 8-Way AMD Opteron High Performance Server

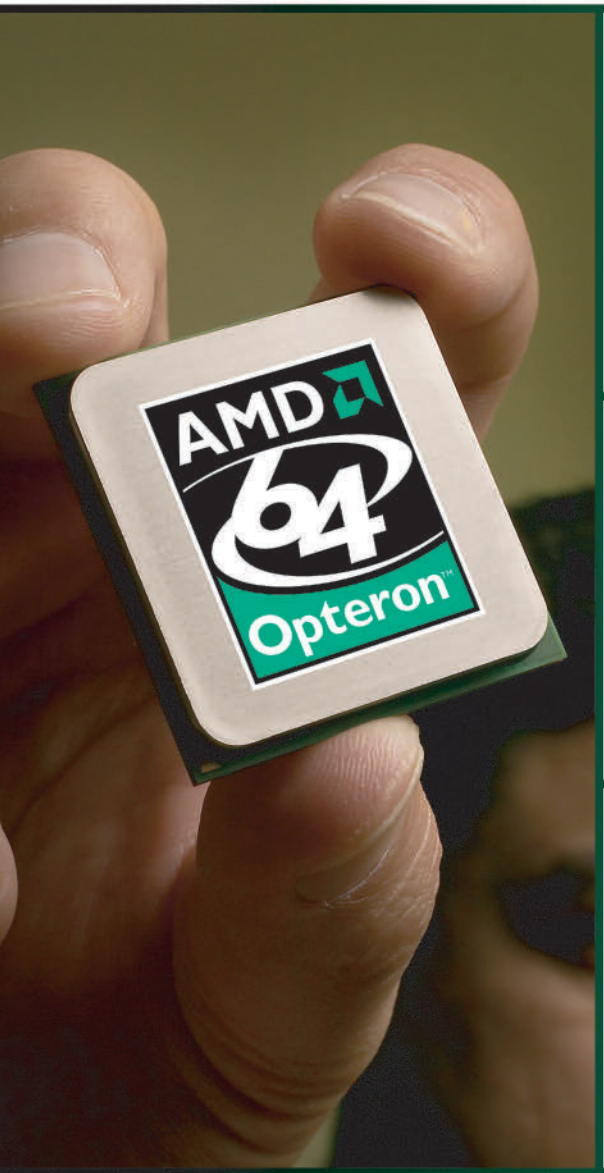

Thinkmate 5U 8-Way Server

- Supports 8-Way AMD Opteron 800 series processors, including Dual-Core.
- Supports up to 128Gb DDR ECC/Registered Memory
- HTX Pro Technology
- 4x 10/100/1000 Ports, 1x 10/100 Port
- 4x PCI-X Expansion Slots

Thinkmate 4U 4-Way Server

- Supports 4-Way AMD Opteron 800 series processors, including Dual-Core.
- Supports up to 64Gb DDR ECC/Registered Memory
- 8x Hotswap Serial-ATA or SCSI Drives
- Up to 3.2 Terabytes of Storage
- 2x 10/100/1000 Ports
- 4x PCI-X Expansion Slots

Thinkmate 1U 4-Way Server

- Supports 4-Way AMD Opteron 800 series processors, including Dual-Core.
- Supports up to 64Gb DDR ECC/Registered Memory
- 2x 10/100/1000 Ports
- Serial-ATA or SCSI Drive
- Exceptional Density

THINKMATE

(800) 371-1212

www.Thinkmate.com