QEMU | BLUETOOTH | SCTP | NAVICORE AND MAEMO MAPPER

LINUX JOURNAL

Since 1994: The Original Magazine of the Linux Community
SEPTEMBER 2007 | ISSUE 161

» Bluetooth Phone Access

» Fedora Directory Server

» Four Powerful Servers

» Stream Control Transmission Protocol

ULTIMATE LINUX HANDHELD

The Nokia N800 is all that and more

+

ULTIMATE LINUX LAPTOP

The Raven X60 Tablet Ultimate Laptop

ULTIMATE LINUX BOX

DIY Ultimate Linux Box for less than $4,0000

www.linuxjournal.com

$5.99US $6.99CAN
09>
0 74470 03102 4

APPRO
HPC Cluster Solutions

Enterprise and High-Performance Computing Under Your Control

Appro is leading 4P x86 computing by combining the latest technology that meets the demands of the enterprise HPC market.

4-Way *Xtreme*Workstation™

- AMD Opteron™ 8000 Series processors
- Up to 128GB of DDR2 533/667 memory
- Up to 6.0TB SATA or 2.4TB SAS
- Hot-swappable drives
- 2 PCI-E x16 slots for high-end graphics card
- Windows® or Linux OS

4-Way 3U *Xtreme*Server™

- AMD Opteron™ 8000 Series processors
- Up to 128GB of DDR2 533/667 memory
- Up to 4.5TB SATA or 1.8TB SAS
- 2 PCI-E x16 and 3 PCI-X slots
- Hot-swappable drives
- Redundant power supplies & fans
- ServerDome Management – IPMI 2.0
- Windows® or Linux OS

For more information, please visit www.appro.com
or call Appro Sales at 800-927-5464

AMD Opteron™ Processors:

- Quad-Core Ready - increase capacity without altering datacenter infrastructure
- Best performance per-watt with energy-efficient DDR2
- Optimized system performance with Direct Connect Architecture

AMD 64 Opteron™

Manage Any Data Center.
Anytime. Anywhere.

Avocent builds hardware and software to access, manage and control any IT asset in your data center, online or offline, keeping it, and your business, "always on".

Visit us on our Remote Control Tour. For locations near you, go to **www.avocent.com/remotecontrol.**

Avocent®
The Power of Being There®

Avocent, the Avocent logo and The Power of Being There, are registered trademarks of Avocent Corporation. ©2007 Avocent Corporation.

CONTENTS

SEPTEMBER 2007
Issue 161

FEATURES

ULTIMATE LINUX BOXES

51 THE ULTIMATE LINUX HANDHELD
Much more than a successor to the Nokia 770.
Doc Searls and Jim Thompson

54 THE ULTIMATE LINUX LAPTOP
256 levels of pressure for this Ultimate Laptop Tablet.
James Gray

58 THE ULTIMATE LINUX BOX
DIY options for the Ultimate or Penultimate Linux Box.
Nicholas Petreley

ON THE COVER
- Bluetooth Phone Access, p. 66
- Fedora Directory Server, p. 72
- Four Powerful Servers, p. 80
- Stream Control Transmission Protocol, p. 88
- Ultimate Linux Handheld, p. 51
- Ultimate Linux Laptop, p. 54
- Ultimate Linux Box, p. 58

Agility.

With Coyote Point, you'll never have to wonder if your network is fast enough or flexible enough. You'll know it is.

From local to global load balancing, application acceleration or ultimate network manageability, Coyote Point leads the pack. We take the guesswork and difficulty out of application traffic management. You won't find anything faster, smarter or more affordable.

Find out why more than 2,000 businesses rely on us to maximize their infrastructure. Learn what Coyote Point could mean to your Web and network traffic.

Write info@coyotepoint.com or call 1-877-367-2696.

Coyotes respond with lightning speed. Your network should, too.

Copyright © 2007 All rights reserved. www.coyotepoint.com

CONTENTS SEPTEMBER 2007 Issue 161

COLUMNS

18 REUVEN M. LERNER'S AT THE FORGE
Database Modeling with Django

26 MARCEL GAGNÉ'S COOKING WITH LINUX
Still Searching for the Ultimate Linux Distro?

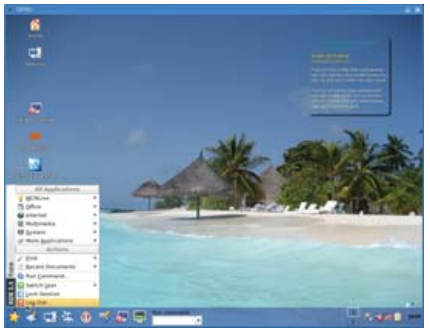

32 DAVE TAYLOR'S WORK THE SHELL
Baccarat Punto Banco, Part II

36 JON "MADDOG" HALL'S BEACHHEAD
Education

38 DOC SEARLS LINUX FOR SUITS
Navigating with the Nokia N800

96 NICHOLAS PETRELEY'S /VAR/OPINION
The Ultimate Linux PVR

IN EVERY ISSUE

8 LETTERS
12 UPFRONT
16 TECH TIPS
48 NEW PRODUCTS
81 ADVERTISERS INDEX

INDEPTH

66 HACKING CELL PHONES VIA BLUETOOTH TOOLS UNDER LINUX
Want to exchange files between PC and cell phone?
Patrick Davila

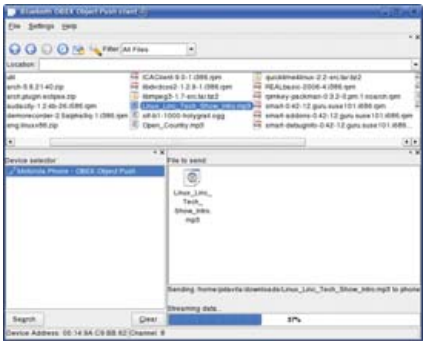

72 FEDORA DIRECTORY SERVER: THE EVOLUTION OF LINUX AUTHENTICATION
Want an alternative to OpenLDAP?
Jeramiah Bowling

80 A $7,000 SERVER COMPARISON
Go big time with your server choice.
Peter Arremann

88 INTRODUCTION TO STREAM CONTROL TRANSMISSION PROTOCOL
Blessed by the IETF.
Jan Newmarch

14 CHUMBY

Next Month

MULTIMEDIA

It's a telephone. No, it's a TV. No, it's a radio. No, it's a calculator. No, it's a video conferencing device. No, the Tornado M20 VoIP phone and digital media center is all of these things and more. It is a remarkable device no bigger than a business phone, and it's all powered by, you guessed it, Linux. Read all about it in next month's issue. And, while we're on the subject of hot devices, we'll get you started programming for the Trolltech Greenphone too.

As always, there's much more. We'll take you on a tour of a Linux-based home management system and show you how nicely the video editor KDENLIVE is coming along. Also in next issue, IBM's Bob Sutor shares his thoughts on open standards and open source.

USPS *LINUX JOURNAL* (ISSN 1075-3583) is published monthly by Belltown Media, Inc., 2211 Norfolk, Ste 514, Houston, TX 77098 USA. Periodicals postage paid at Houston, Texas and at additional mailing offices. Cover price is $5.99 US. Subscription rate is $25/year in the United States, $32 in Canada and Mexico, $62 elsewhere. POSTMASTER: Please send address changes to *Linux Journal*, PO Box 980985, Houston, TX 77098. Subscriptions start with the next issue.

Don't feel bad, Our servers won't go down on you either.

We've all known disappointment. And few things are more disappointing than undependable, expensive servers that don't satisfy your needs. To avoid the heartbreak of lost expectations, QSOL provides dependable, server solutions for virtually any popular operating system at a price that will blow your mind not your budget.

QSOL.com Server Appliances come in a variety of configurations to meet the most specific needs of our users. Whether you're into 1U with 8/16 Cores, or 5U with nearly 40+ Terabytes, we can provide a server that can perform under the most demanding conditions. If your server isn't giving you what you want, visit QSOL. (Don't worry, it's a sure thing!)

QSOL.COM
Server Appliances

Serving Servers since 1999

© 2007 QSOL.COM INC. All rights reserved, all trademarks are the property of their respective companies.
The person pictured in this advertisement is a model and is used for illustrative purposes only.

LINUX JOURNAL™
Since 1994: The Original Magazine of the Linux Community

Digital Edition Now Available!

Read it first
Get the latest issue before it hits the newsstand

Keyword searchable
Find a topic or name in seconds

Paperless archives
Download to your computer for convenient offline reading

Same great magazine
Read each issue in high-quality PDF

Try a Sample Issue!
www.linuxjournal.com/digital

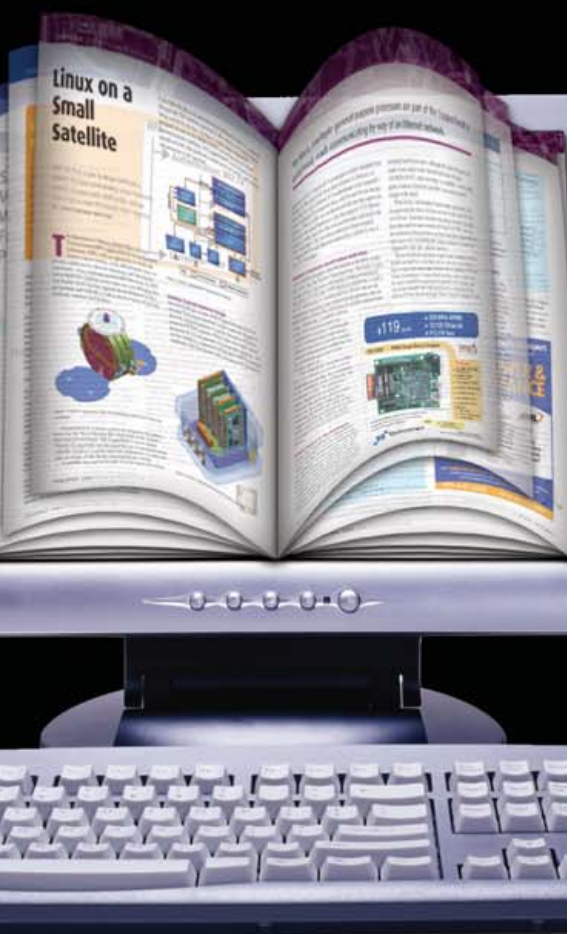

LINUX JOURNAL

Editor in Chief
Nick Petreley, ljeditor@linuxjournal.com

Executive Editor	Jill Franklin jill@linuxjournal.com
Senior Editor	Doc Searls doc@linuxjournal.com
Art Director	Garrick Antikajian garrick@linuxjournal.com
Products Editor	James Gray newproducts@linuxjournal.com
Editor Emeritus	Don Marti dmarti@linuxjournal.com
Technical Editor	Michael Baxter mab@cruzio.com
Senior Columnist	Reuven Lerner reuven@lerner.co.il
Chef Français	Marcel Gagné mggagne@salmar.com
Security Editor	Mick Bauer mick@visi.com

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

Proofreader Geri Gale

Publisher	Carlie Fairchild publisher@linuxjournal.com
General Manager	Rebecca Cassity rebecca@linuxjournal.com
Director of Sales	Laura Whiteman laura@linuxjournal.com
Regional Sales Manager	Joseph Krack joseph@linuxjournal.com
Regional Sales Manager	Kathleen Boyle kathleen@linuxjournal.com
Circulation Director	Mark Irgang mark@linuxjournal.com
Marketing Coordinator	Lana Newlander mktg@linuxjournal.com
System Administrator	Mitch Frazier sysadm@linuxjournal.com
Webmaster	Keith Daniels webmaster@linuxjournal.com
Accountant	Candy Beauchamp acct@linuxjournal.com

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

Editorial Advisory 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

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

Subscriptions
E-MAIL: subs@linuxjournal.com
URL: www.linuxjournal.com/subscribe
PHONE: +1 713-589-3503
FAX: +1 713-589-2677
TOLL-FREE: 1-888-66-LINUX
MAIL: PO Box 980985, Houston, TX 77098 USA
Please allow 4–6 weeks for processing address changes and orders
PRINTED IN USA

LINUX is a registered trademark of Linus Torvalds.

Power of 32 Processors

AMD 64 Opteron

256GB of Memory

Based on industry leading AMD Opteron™ microprocessors, the **HPC A5808-32** server provides industry leading scalable x64 processing with the latest I/O, network, memory, and power efficient technologies.

With features such as **AMD quad-core Opteron™** microprocessors and up to 256GB of DDR2 for memory-intensive performance, multithreaded applications like CRM, ERP, e-commerce, and virtualization will see significant performance improvements.

And **HPC MasterSight™** delivers features to help manage the system with inclusive diagnostic tools.

Specializing in High Performance Computers, HPC systems, Inc. is a solution provider serving the financial, health, educational, and government sector. We appreciate that you have a choice of dozens of vendors, but not all of them have our uncompromising dedication and experience in producing the best solutions.

HPC SYSTEMS INCORPORATED

HPC Systems, Inc.
48009 Fremont Blvd
Fremont, CA 94538
Tel: 510-656-8282
Fax: 510-656-8341
E-mail: info@hpcsystems.com

AMD Solution Provider PLATINUM

SBA 8(A) CERTIFIED

GSA GS-35F-0596R Contract Holder

AMD, AMD Opteron, and the AMD logo are registered trademarks of the Advanced Micro Devices or its subsidiaries in the US and other countries.

letters

And Now for Something Completely Different

There is a nice Python tutorial in the June 2007 issue ["Programming Python, Part I" by José P. E. Fernandez]. It would be wonderful if, after the tutorial is done, a monthly Python column would appear in the pages of *Linux Journal*. (I have seen columns on other languages, for example, Perl, but never a column on Python.)

--

Richard

BSD Script Modification

In the "Displaying Image Directories in Apache, Part IV" article by Dave Taylor [July 2007], I ran into a problem using the script on my hosting company's BSD-based system. Specifically, the problem was in the figuresize function returning an invalid width value.

The figuresize function can be changed to the following:

```
figuresize()
{
    width="$(identify -format %w $1)"
   height="$(identify -format %h $1)"
}
```

This change solved my problem and makes the function more efficient by eliminating four invocations of cut for every image.

--

Doug Winterburn

Recognize This

I've been more than getting my money's worth from *LJ* just from Dave Taylor's Work the Shell articles, but this month, my introduction to Tesseract doubled my pleasure ["Tesseract: an Open-Source Optical Character Recognition Engine" by Anthony Kay, July 2007]. I have been looking for an OCR program since I quit Microsoft seven years ago, and now I have one. Tesseract is outstanding, and Anthony Kay did a great job with the article.

--

Bruce Bales

Error Handling Instead of Ignoring

I am writing regarding to the article "Writing Your Own Image Gallery Application with the UNIX Shell" by Girish Venkatachalam, published in the July 2007 issue of *Linux Journal*.

In the script on page 71, Girish suppressed the mkdir error by redirecting the error message to the bit bucket:

```
# we don't want mkdir shouting at us for
# directories that exist!
mkdir $dimension 2>/dev/null
```

My suggestion is to do it this way instead:

```
# use -e instead of -d, since an existing
# file with the same name could also prevent
# you from creating the directory
[ -e $dimension ] || mkdir $dimension
```

Even better, handle the mkdir error this way after the mkdir error:

```
if [ $? -ne 0 ]; then
        echo Handle my error here
fi
```

Thanks for the great magazine, and keep up the good work.

--

Jack

Response to the "Don't Just Beat Me, Teach Me" Letter, July 2007

Writing instructional software from scratch for Linux (or any other OS for that matter) is a time-consuming and nontrivial activity. For people to invest their time and energy in this, there would definitely need to be a payoff to make it a worthwhile undertaking. I think a much more promising approach would be a Linux port of existing instructional software.

ChessMaster is indeed a very fine instructional package, and I highly recommend it to my students. Several years ago, I contacted Ubisoft regarding a Linux version of ChessMaster. Unfortunately, the predictable response was "not planned for the foreseeable future". Linux has gained a lot of traction since then, even in relatively small market niches like chess.

A good example is one of the world's strongest programs, Shredder (**www.shredderchess.com**), which has been made available for Linux. I also happen to work as a consultant for a North American distributor and retailer (**www.chesscountry.com**) of chess software. My recommendation for Convekta (which produces very good instructional software) Linux software was favorably received, although I can't make any predictions.

As Linux gains critical mass in the chess sphere, companies like Ubisoft will find it difficult to ignore Linux lest they risk being displaced by a newcomer. An intense lobbying effort might be persuasive, and they very well might consider porting to Linux. There certainly are enough capable developers available to make this happen.

Until then, I unfortunately have no better recommendation than to make do with one of the Windows-Linux integration techniques that are available to us. Mr Colon is quite right not to struggle with Wine. I have made the determination that running Windows apps via Wine is hit or miss—half will run after intense configuration effort, and half won't run at all.

The two preferable options are VMware or VNC. If you have only one computer and money is no object, VMware might be the way to go. My preferred solution is VNC (which is free). I utilize a multiboot laptop for my IT-consulting work. The laptop can

servers DIRECT™

MORE PRODUCTS, BETTER SERVICE, GUARANTEED.

GO STRAIGHT TO THE SOURCE!

KEEP YOUR BUSINESS RUNNING SMOOTHLY

PROTECT YOUR SMALL BUSINESS WITH THE BUILT-IN SECURITY ENHANCEMENTS OF THE DUAL-CORE INTEL® XEON® PROCESSOR IN YOUR SERVERSDIRECT SYSTEM.

SDR-1105T 1U ENTRY LEVEL SERVER

$959.99

EXCELLENT GENERAL PURPOSE SERVER FOR ORGANIZATIONS WITH THE NEED FOR A LOW, ENTRY LEVEL PRICE

* 1U Rackmount Chassis with 520W power supply
* Supermicro X7DVL-L Server Board with Intel® 5000V (Blackford VS) Chipset
* Intel® Dual-Core Xeon Processor 5050 3.0GHZ 667 MHz
* Total 512MB, 2pcs x 256MB Kingston DDR2 533Mhz FB-DIMM ECC
* Seagate SATAII 80GB 7200 RPM 8MB Cache SATA 3.0Gb/s Hard Drive
* 4 x 1" Hot-swap SATA Drive Bays
* Two Intel® 82563EB Dual-port Gigabit Ethernet Controller
* Intel® ESB2 SATA 3.0Gbps Controller RAID 0, 1, 5, 10 support

STARTING PRICE $1,199

SDR-2503T 2U APPLICATION SERVER

Highest performing with Dual Core/ Quad Core Xeon CPU based. Excellent with general purpose applications and provide the most power.

* 2U Rackmount Chassis with 650W power supply
* Supermicro X7DVL-E Server Board with Intel® 5000V (Blackford VS) Chipset
* Intel® Dual-Core Xeon Processor 5050 3.0GHZ 667 MHz
* Total 512MB, 2pcs x 256MB Kingston DDR2 533Mhz FB-DIMM ECC
* Seagate SATAII 250GB 7200 RPM 8MB Cache SATA 3.0Gb/s Hard Drive
* 6 x 1" Hot-swap SATA Drive Bays
* Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
* Intel® ESB2 SATA 3.0Gbps Controller RAID 0, 1, 5, 10 support

STARTING PRICE $2,199

SDR-3500T 3U DATABASE SERVER

Easily Scalable storage solution with hot-swap functionality for growing businesses

* 3U Rackmount chassis with Redundant 800W power supply
* Supermicro X7DBE+ Server Board with Intel® 5000P (Blackford) Chipset
* Intel Quad-Core Xeon Processor E5310 1.6GHZ
* Total 1024MB, 2pcs x 512MB Kingston DDR2 533MHz FB-DIMM ECC
* Seagate SATAII 500GB 7200 RPM 16MB Cache SATA 3.0Gb/s Hard Drive
* 16 x 1" Hot-swap SATA Drive Bays
* Dual-port Gigabit Ethernet Controller
* Intel ESB2 SATA 3.0Gbps Controller RAID 0, 1, 10 support

STARTING PRICE $1,999

SDR-7045B-TB 4U FILE SERVER

4U Quad-Core Xeon Server offers excellent value and expandability

* 4U Rackmountable / Tower with 650W power supply
* Supermicro Super X7DBE Server Board with Intel® 5000P (Blackford) Chipset
* Intel Quad-Core Xeon Processor E5310 1.6GHZ
* Total 1024MB, 2pcs x 512MB Kingston DDR2 667MHz FB-DIMM ECC
* Seagate SATAII 750GB 7200 RPM 16MB Cache SATA 3.0Gb/s Hard Drive
* 6 x 3.5" Hot-swap SAS/SATA Drive Bays
* Dual-port Gigabit Ethernet Controller
* Intel ESB2 SATA 3.0Gbps Controller RAID 0, 1, 5, 10 support

STARTING PRICE $3,299

SDR-5111T 5U ADVANCED STORAGE SERVER

Quad Core dual Xeon CPU based, with 24 hot-swap hard disk bays suitable for 18TB of pure data Storage capacity

* 5U Rackmount chassis with Redundant 1350W power supply
* Supermicro X7DBE Server Board with Intel® 5000P (Blackford) Chipset
* Intel Quad-Core Xeon Processor E5310 1.6GHZ
* Total 1024MB, 2pcs x 512MB Kingston DDR2 667MHz FB-DIMM ECC
* Seagate 750GB 7200 RPM 16MB Cache SATA 3.0Gb/s Hard Drive
* 24 x 1" Hot-swap Drive Bays
* Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
* Intel ESB2 SATA 3.0Gbps Controller RAID 0, 1, 5, 10 support

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

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

1.877.727.7887 | www.ServersDirect.com

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

intel Xeon® inside™

Dual-core. Do more.

boot Debian, Red Hat or XP (I recently removed SUSE for reasons covered in recent issues of *LJ*). If I need to access XP from my desktop, I simply connect the laptop to my home wireless network, boot XP and start a VNC server. Windows shares are made visible via Samba.

The combination of VNC and Samba yields complete seamless access to the laptop running XP from my Debian workstation. Despite the fact that all my hosts are on a wireless network, the VNC performance is actually quite good. Of course, if you have an exclusively wired network, performance will even be better.

I think the VNC solution is a good fit for most Linux users, as most tend to have more than one computer that already may be networked together. Adding network connectivity is relatively inexpensive as compared to going the VMware route.

Unfortunately, I don't know of a Linux equivalent for ChessMaster, but until one arrives on the scene, I hope these suggestions will tide ChessMaster users over.

--

Peter Stein

Picture Imperfect

I just noticed the Pixel article ["Interview with Pavel Kanzelsberger, Creator of Pixel" by James Gray, July 2007]. Please do not promote that project; it is nothing more than a fraud.

You should really look through their forums. There are many people who have been scammed. It is not current, it does not work, and there is a one-man team working on it. Besides that, it is closed source. The project is doomed and a scam, and I just ask you to please let it die.

--

Anonymous

256MB of RAM Is Plenty

In Nicholas Petreley's "Amazing Free Distributions Abound" [July 2007], it says, "I run Damn Small Linux on my old Compaq notebook with 256MB of RAM simply because it won't run anything more bloated than that", but I run a full Ubuntu install on a Pentium 2 with 192MB of RAM. It's a bit slow, but not as bad as Windows XP was.

I'm going to install XFCE instead of GNOME to give it a little boost, but 256MB is plenty of RAM for the average distro. Debian + IceWM would be pretty fast too.

--

Mackenzie Morgan

Let's Make a Deal

Nicholas, Your distro shopping spree in July's issue was entertaining, but it would be even nicer to read about other source-based distros besides Gentoo. Any chance of that happening?

I'm probably not the first one to bring this up, but I have to ask you if your anticipation of the next Linspire has faded by now. The recent deal between Microsoft and Linspire must have been announced right when you guys were busy getting the issue printed.

--

Juuso

I'm very disappointed in the recent deals between Microsoft and Linspire, Xandros and others. Deals like these with Microsoft usually send Microsoft's new "partners" to the morgue. Time will tell if Linspire and others find a way to take the money and run or if Microsoft has the last laugh.—Ed.

LINUX JOURNAL

At Your Service

MAGAZINE

PRINT SUBSCRIPTIONS: Renewing your subscription, changing your address, paying your invoice, viewing your account details or other subscription inquiries can instantly be done on-line, **www.linuxjournal.com/subs**. Alternatively, within the U.S. and Canada, you may call us toll-free 1-888-66-LINUX (54689), or internationally +1-713-589-2677. E-mail us at subs@linuxjournal.com or reach us via postal mail, Linux Journal, PO Box 980985, Houston, TX 77098-0985 USA. Please remember to include your complete name and address when contacting us.

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

LETTERS TO THE EDITOR: We welcome your letters and encourage you to submit them to ljeditor@linuxjournal.com or mail them to Linux Journal, 1752 NW Market Street, #200, Seattle, WA 98107 USA. Letters may be edited for space and clarity.

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

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

ON-LINE

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

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

Maximize **AMD Opteron™ Quad-Core** Performance with **UniServer**

Up to 32GB memory on 8 DIMM sockets

1U UniServer with 1 Socket

- Dual/Quad-Core AMD Opteron™ 2000 series
- Up to 32GB of memory
- 2 hot swap SATA 3.5" HDDs
- MCP55V Pro chipset
- Slim CD/DVD-ROM
- 1 PCI-Express x16 slot
- Remote management-IPMI 2.0

Find Full line of Innovative AMD Opteron™ Servers and Workstation Platforms

1U Server: 1 & 2 Sockets, up to 64GB memory
2U Server: 2 & 4 Sockets, up to 128GB memory
3U Server: 2 & 4 Sockets, up to 128GB memory
Workstation: 1, 2 & 4 Sockets, up to 128GB memory

BELL MICRO www.bellmicro.com 1-800-291-2070

AVNET Applied Computing Solutions www.avnet.com 1-888-300-8277

SYNNEX CORPORATION www.synnex.com 1-888-756-4888

UNIWIDE Technologies

www.**uniwide**.com

1-877-520-0071

Uniwide is an official AMD Validated Server Program Partner

Copyright © 2007 Uniwide Technologies, Inc. All right reserved. Specifications are subject to change without notice.

AMD 64 Opteron™

diff -u

WHAT'S NEW IN KERNEL DEVELOPMENT

UMSDOS, once a proud gateway into the Linux world for MS-DOS users, has been gone from the kernel for a while, and now even its dregs are being culled, bit by bit. **Jesper Juhl** has posted patches to remove the configuration entries from kconfig, as well as the ioctls UMSDOS had used. But, it turns out we still can't get rid of the ioctls, because the software of the future must know not to reuse the numbers. So, like DevFS, it may be impossible to remove all evidence of UMSDOS fully.

Intel has released **PowerTOP**, a standalone utility to help identify how much power is being used by the various applications and other parts of the system, including kernel drivers. The goal is to help all software developers trim down their power usage to extend laptop uptime greatly. **Arjan van de Ven**, who announced the project on the linux-kernel mailing list, said that on a test laptop Intel had been able to increase battery life by more than an hour.

Dave Jones has posted a patch to undeprecate the **Raw I/O driver**. This driver is no longer needed because users simply can open the target device with the O_Direct flag. But evidently, there is a significant number of users who are just not in a position to change their software. So, even though a better method exists, it looks like the Raw driver will be sticking around for the foreseeable future. Even listing it as obsolete is now a no-no, because, as Dave says, it would just lead someone to add the driver to the "to be removed" list, which would start up the whole discussion again.

Geert Uytterhoeven posted patches for a **Flash ROM storage driver** for the **PS3**. This seemed to be a welcome set of patches, with no significant technical obstacles.

Jesse Barnes at **Intel**, working with the **Framebuffer** developers, is attempting to change the way graphics are handled in the Linux world completely. They want to enhance the kernel's graphics support to enable full-featured non-X Window System graphics capabilities to Linux. As it turns out, this is a very controversial project, and it represents another in a long line of attempts to replace the X Window System. The problems with such a project are enormous, not the least of which will be that any alternative to the X Window System will have to provide a competitive set of features, making it at least as large as the X Window System itself. This prospect will never appeal to kernel developers who feel that a small, sleek kernel is best. But, it's also true that graphics devices have been a strange exception to the general rule that hardware controllers belong in the kernel. It'll be very interesting to see how this particular attempt at a kernel graphics system plays out.

Rob Landley has decided to reorganize a bit of the **Documentation directory**, creating a directory for the **Amiga architecture**, and putting arm, cris, blackfin, parisc, powerpc, s390, x86_64 and uml into that directory. While he was at it, he moved the Multiport Serio IO cards into the serial directory.

—ZACK BROWN

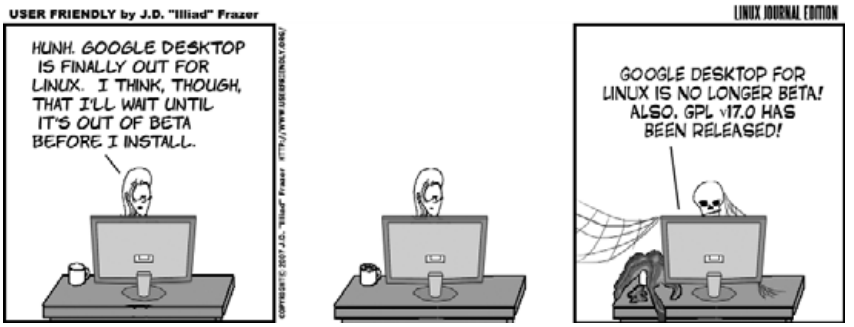

LJ Index, September 2007

1. Dollars spent by Nike on its swoosh logo: 35
2. Billions of dollars spent on marketing last year by Nike: 1.7
3. Percentage of Web sites that are pornographic: 12
4. Pornographic percentage of all search engine requests: 25
5. Porn percentage of all downloads: 35
6. Number of users viewing porn per second: 28,358
7. Dollars (US) spent on port every second: 89
8. Number of new porn sites going up every day: 266
9. Estimated billions of porn Web pages: .372
10. Position of "sex" among the most searched words: 1
11. US revenue from Internet porn in 2006, in billions of dollars: 2.84
12. Male percentage of Internet porn users: 72
13. Percentage of porn traffic during the 9–5 eight-hour workday: 70
14. Percentage of porn sites produced by the US: 89
15. Position of AdultFriendFinder among most popular porn sites: 1
16. Number of AdultFriendFinder sites followed by Netcraft: 75
17. Number of AdultFriendFinder sites known to be served by Unknown: 46
18. Number of AdultFriendFinder sites known to be served by Windows: 1
19. Number of AdultFriendFinder sites known to be served by BSD: 2
20. Number of AdultFriendFinder sites known to be served by Linux: 28

Sources: 1–15: *Good Magazine*
16–20: Netcraft.com

—Doc Searls

SMART, CURIOUS, INQUISITIVE, MOTIVATED. Daring to take on the challenges that others won't. This is a White Oak employee on any given day.

Tackling ambitious objectives that push your limits. Working with passionate people who inspire greatness. Finding ingenious solutions to the toughest problems.

On the job or off, White Oak people keep their adrenaline flowing!

Our benefits are as exceptional as the people we hire. If you have impressive architecture and development experience in one or more of the following areas, show us what you can do. Send an email to: bruce@woti.com.

- ***Web and Database Technologies***
- ***C++, Python, Perl , Linux***
- ***User Interface Development***
- ***Parser & Compiler Design***
- ***TCP/IP Programming***

Positions located in the Washington, D.C. area.

www.woti.com

US Citizenship required. White Oak Technologies, Inc. is an Equal Employment Opportunity Employer.

Chumby rasa

The question at this point is whether the name Chumby will come to stand for every other hackable pillow-like device. (Such was the fate of Kleenex and Escalator.) If all goes according to plan, Chumbys should be on the market by now. Prototypes and development versions have been circulating for about a year, and a sizable development community has grown around it. Given how much it has grown and changed in the public womb, there's no telling how it'll evolve out in meatspace.

"Chumby Industries was formed by hackers who wanted to create something interesting, useful and different. The starting point was the humble clock radio", its creators explain. Since then, Chumby has evolved from a clock in a cushion to an Any-purpose Net-native Linux device. That's any with a capital A, because the Chumby is built to be hackable at every level, including the physical. Not only does it sense motions and squeezings, but it also hosts an assortment of charms, through its "outerware API". The charms and much more about the Chumby were designed by Susan Kare (who designed the original desktop icons for the Macintosh and Windows, among too many other things to mention). Susan is Creative Director for Chumby. The company might be cuddly, but it means business too.

Photo by Bunnie Huang, VP Hardware Engineering and Founder of Chumby

Tech details: 3.5" LCD color touchscreen, two external USB 2.0 full-speed ports, 350MHz ARM controller, 64MB SDRAM, 64MB NAND Flash ROM, 2-Watt amp with stereo speakers, headphone output, squeeze sensor, accelerometer, Wi-Fi connectivity and microphone—plus whatever you can do with it. For more information, visit **www.chumby.com**.

—DOC SEARLS

The World Continues to Await Dell Linux

Dell's IdeaStorm site opened earlier this year to a hurricane of strong input from the Linux community. As we reported in the June 2007 issue, nearly all of the top ten vote-getters among new ideas for the company were Linux- or open-source-related. Since then, Dell has committed to making Linux-based desktops and laptops, and it offers a nice array of product choices based on Red Hat, SUSE and Ubuntu distros—in the US, that is.

The top topic on the IdeaStorm site as of mid-June 2007 was "Sell Linux PCs Worldwide—not only the United States". It has more than 24,000 votes. Next in line was "Dell Ubuntu for Europe", with 11,370 votes. Other items on the same page were "Ubuntu Dell Must Cost Le$$ Than Windows Dell", "Same discounts available on Ubuntu and Windows", "Provide Linux Drivers for all your Hardware", "Pre-Installed OpenOffice.org | alternative to MS Works & MS Office", "Have Firefox pre-installed as default browser" and "TV Commercial for New Ubuntu PCs".

For its part, Dell says it is simply starting in the US and is "still working out details of its global programme". Support in languages other than English is an issue. Look up "linux" at **www.google.com/trends**, and you'll find the top country sources of Linux searches are Czech Republic, Russia, Bulgaria, Indonesia, India, Slovakia, Romania, Ukraine, Hungary and Poland. You'll find similar results with searches for any of the distros. The US isn't in the top ten for any of them.

There's also room for improvement in the US. Dell still buries Linux details deep in its Web site. Although plenty of pages are devoted to Linux or Linux-related products, the word "linux" cannot be found on the home page or any pages up to two layers down the Dell site schema—not even in the pages for enterprise servers (as it is, say, for HP). Yet, there are plenty of "Dell recommends Windows Vista—Home Premium" messages at the tops of Dell product pages.

Meanwhile, Michael Dell's interest doesn't seem to be flagging. As of press time, the top computer on Michael's list is still a Dell Precision M90 running Ubuntu 7.04 Feisty Fawn (**www.dell.com/content/topics/global.aspx/corp/biographies/en/msd_computers?c=us&l=en&s=corp**).

—DOC SEARLS

They Said It

There's really only one rule for community as far as I'm concerned, and it's this—in order to call some gathering of people a "community", it is a requirement that if you're a member of the community, and one day you stop showing up, people will come looking for you to see where you went.
—Adam Fields, www.aquick.org/blog/2007/05/15/the-first-rule-of-community

Every Web service has ended up being its own little snowflake, needing its own special treatment. The thing is, it's better than nothing so we trudge along and build the client libraries anyway.
—Les Orchard, blogs.opml.org/decafbad/2007/05/27#When:2:35:59PM

Les Orchard says each API is a snowflake that every developer has to build custom interfaces for. Someday all these guys will realize they need to get together and come up with some standards for serializing data to simplify the work for themselves and developers. A lot of compromise and working together will be needed to make this happen. And when they have finished, many years from now, they will be where we were with XML-RPC in 1998.
—Dave Winer, www.scripting.com/stories/2007/05/27/linksOnlyANerdCouldLove.html

A complex system that works is invariably found to have evolved from a simple system that works.
—John Gall's 15th law of Semantics, en.wikipedia.org/wiki/Systemantics

It's reasonable to talk about software as being alive. It's symbiotic. It needs a host geek in which to live.
—Jeremy Ruston (author of TiddlyWiki), from a talk about TiddlyWiki in London

Looking for hardware as killer as your next app?

GEMINI - Intel® Core™ 2 Duo

SBC-GX533 - AMD Geode™ GX

ZEUS - PXA270

VIPER - PXA255

VULCAN - IXP425

...look to Arcom embedded boards

- Intel XScale, Core 2 Duo, AMD and X86 technology
- Fanless solutions
- Display controller for TFT and STN displays
- Serial ports for legacy communications
- USB, SDIO, CompactFlash, PC/104 & PCI Expansion
- Complete operating system and application development environment

Microsoft® Windows® Embedded Partner
Gold-level Member

Linux 2.6

IBM WebSphere® Everyplace Micro Environment

Java POWERED

eclipse

888-941-2224
www.arcom.com

Arcom
A MEMBER OF EUROTECH GROUP

TECH TIPS

Automount FTP sites as filesystems and take the easy road to installing graphics drivers.

Combine Automount with FUSE and CurlFtpFs

Some time ago, I "discovered" automount, and after using it for a while, I wondered if it would be possible to combine it with FUSE and CurlFtpFs to make automatic filesystem access possible to FTP sites. This wasn't very trivial, because the automount software had some problems with the interpretation of the map file. I solved this problem by creating a helper script faking a curl filesystem.

Here are the steps:

1) If not yet enabled in the kernel, enable autofs by setting CONFIG_AUTOFS and CONFIG_AUTOFS4 to yes or module and rebuild the kernel.

2) Get FUSE from **fuse.sourceforge.net** and install it.

3) Get CurlFtpFs from **curlftpfs.sourceforge.net** and install it.

4) Create the map file /etc/auto.ftp:

```
-fstype=curl,allow_other,ro        :ftp\://&/
```

This will tell the automounter to use the curl filesystem to mount an FTP site. I added the allow_other option so anyone on the system can use this method. See the documentation for FUSE and CurlFtpFs for other possible options.

5) I created the following helper script, /sbin/mount.curl:

```
#! /bin/sh

mount -t fuse curlftpfs
```

This will be used by mount to mount the curl filesystem, but it effectively uses the FUSE filesystem with CurlFtpFs to mount the FTP site.

6) Create the directory /mnt/ftp (or any other you like).

7) Then, after issuing the command (as root):

```
automount /mnt/ftp file /etc/auto.ftp
```

you can access FTP archives simply by changing to the directory /mnt/ftp/ftp.linuxjournal.com, as if on your own computer. After some time, automount (the default is five minutes, but that can be changed with the -t option with automount) will unmount this directory again and release the connection to the FTP site. Don't forget that every time you access a file, it will be transferred to you via FTP, which can take some time depending on your Internet speed.

—Michiel, from somewhere in cyberspace.

Easy Installation of NVIDIA and ATI Drivers on Ubuntu/Kubuntu

Check the Ultimate Linux Box article in this issue (page 64), and you'll find the "hard way" instructions for installing the latest NVIDIA driver on Ubuntu/Kubuntu. There's an easier way to install NVIDIA (or even ATI) drivers on Ubuntu/Kubuntu and all its spin-offs. You can use a program called Envy, which is designed to automate installation of accelerated graphics drivers.

I heard about Envy but opted to go the "hard way" at first because I'd read numerous reports of Envy failing to work. The hard way isn't very hard for me anymore, because I've gone through the process so many times. It's a shame I wasted my time though, because I found out later that Envy is, in fact, very easy. And, it works just fine with the latest Ubuntu/Kubuntu, Linux Mint and other Ubuntu spin-off distros.

Figure 1. You can run Envy from a console or terminal window.

Figure 2. You can run a graphical version of Envy if you managed to get to a GUI desktop.

You can use a program called Envy, which is designed to automate installation of accelerated graphics drivers.

The first thing you need to do is install Envy. Run the command:

```
sudo apt-get install envy
```

You may find that it automatically installs a number of dependencies.

Use the command `envy -t` to run Envy with the text-mode interface. This is especially useful if you weren't able to get a graphical desktop running at all, because you can run this from a text console. It works just as well in a terminal window on a graphical desktop though. See Figure 1 for a picture of the text-mode main menu.

You can run a graphical version of Envy instead, with the command `envy -g`. See Figure 2 for a picture of the graphical main menu.

Select the first menu choice for the NVIDIA driver. You'll have to enter your password if you ran Envy as a normal user. Then, follow the prompts. It will ask you if you want to update your /etc/X11/xorg.conf file. The default answer is "y", and I recommend you use it.

If you installed Linux and got a graphical desktop with low resolution because it couldn't detect your graphics card properly, you probably won't want to stick with that low resolution. The envy program won't necessarily correct this problem for you. You need to change your Screen section in the /etc/X11/xorg.conf file. For example, I deleted the resolution on the list starting with 1024x768 and replaced it with a 1920x1200 resolution, the only one I use:

```
Section "Screen"
    Identifier     "Default Screen"
    Device         "nVidia Corporation"
    Monitor        "Generic Monitor"
    DefaultDepth    24
    SubSection     "Display"
        Depth       24
        Modes      "1920x1200"
    EndSubSection
EndSection
```

—Nicholas Petreley

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

When Chris Negus speaks, people learn Linux®!

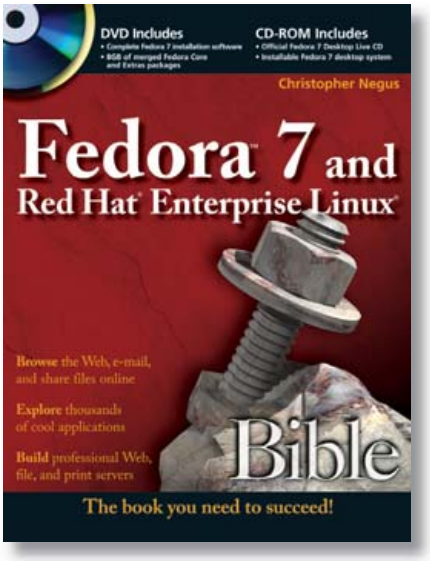

Chris Negus is back with the only book you'll need on Fedora 7 and Red Hat Enterprise Linux—from the basics up to advanced system administration skills. Packed with tutorials, techniques, and the latest on platforms, 3D interfaces, and media, this is the best Fedora book on the market.

978-0-470-13075-9

Available wherever books are sold.

BICENTENNIAL 1807 WILEY 2007 BICENTENNIAL

wiley.com

Wiley and the Wiley logo are registered trademarks of John Wiley & Sons, Inc. All other trademarks are the property of their respective owners.

Database Modeling with Django

REUVEN M. LERNER

Let Django and its object-relational model do the SQL database work for you.

The past few months, this column has been examining Django, a popular open-source Web development framework written in Python. Django sometimes is described as a rival to Ruby on Rails or a Python version of Rails, but just as Python and Ruby are distinct languages, each with its own strengths and weaknesses, Django and Rails are different frameworks, and each has its own set of trade-offs.

If you have been following this series of columns about Django, you already have seen how to download and install the Django software, how to create and configure a site and application, and even how to create views (Python methods that handle the business logic) and templates (HTML files with special rules for interpolating variables and dynamic content). With everything we've looked at so far, you could presumably create an interesting dynamic Web application.

However, most modern Web applications have another component, a relational database, on which they rely for data storage and retrieval. Sure, you could store everything on the filesystem or even in memory, but for most of us, a relational database is the path of least resistance, ensuring the safety of our data while providing a great deal of flexibility in retrieving it.

This month's column, then, looks at the ways in which Django programmers can store and retrieve information in a database. If you have worked with databases only from PHP or CGI programs, you will be surprised and impressed by the degree of automation Django provides. If you have worked with Ruby on Rails, you probably will think the Django programmers are working too hard—to which Django hackers would say that they want to have full control over their application, rather than rely on behind-the-scenes magic.

Creating a Model

The term model in the Django world describes a Python object for which there is a persistent state, presumably stored in a relational database. We don't need to use models to integrate a database into Django, but it would be difficult (not to mention unaesthetic) if we were simply to stick SQL queries into our templates. So instead, we use Django's built-in object-relational mapper, working solely with objects from within our views and templates. The mapper's job is to translate our method calls into SQL and then translate the resulting database response into Python objects.

But, before we even can create our model object, we first must have a database table to which the object will connect. Django requires that we define our model using Python code, describing the table's name, fields and even some default values.

If we were interested in keeping the blog application we started last month, we probably could define our table in PostgreSQL as follows:

```
CREATE TABLE Posting (
id          SERIAL     NOT NULL,
title     TEXT     NOT NULL,
body     TEXT     NOT NULL,
posted_at     TIMESTAMP NOT NULL DEFAULT NOW(),

PRIMARY KEY(id)
);
```

But in Django, we don't create the above SQL directly. Rather, we use Python to create it for us. For example, we can define the above table in Django as follows:

```
from django.db import models

class Posting(models.Model):
    title = models.CharField(maxlength=30)
    body = models.TextField()
    publication_date = models.DateTimeField()
```

Our model is a Python class, which inherits from django.db.models.Model. We define each field with a particular type, using objects that we imported from django.db.models. As shown above, some of these data types can be restricted or modified from their defaults by passing parameters. Some are defined specifically because they have built-in limits, such as EmailField, which must be a valid e-mail address. By defining columns with ManyToManyField and ForeignKey objects, it's possible to define a variety of relationships among tables.

The above code should be placed in models.py, a

Python file that sits within our application's directory (blog in this case), which itself sits inside our Django site directory (mysite in this case). Thus, the models for my blog application reside in mysite/blog/models.py, whereas the models for a poll application would reside in mysite/poll/models.py.

Notice that we don't have to define a primary key, which traditionally is called id and is a nonrepeating integer. (In PostgreSQL, we set it to have a SERIAL data type, which gives the column a default value taken from a newly created sequence object. In MySQL, you would set the column to AUTO_INCREMENT, which has some of the same capabilities as a sequence.) Django creates the id column for us automatically. Django handles potential namespace conflicts by prefacing the table name with the application name. So, the posting table within the blog application becomes the blog_posting table.

Now, how do we turn our Python code into SQL? First, we have to be sure Django knows which database to use. If you have been following along since my first Django article in the July 2007 issue, you already have added the appropriate lines to settings.py, a site-wide configuration file in which we define the database type, name, user and password. Here are the values that I have installed:

```
DATABASE_ENGINE = 'postgresql'
DATABASE_NAME = 'atf'
DATABASE_USER = 'reuven'
DATABASE_PASSWORD = ''
DATABASE_HOST = ''
DATABASE_PORT = '5433'
```

It's also important to check that the application is defined in INSTALLED_APPS, a tuple of strings. On my system, INSTALLED_APPS looks like this:

```
INSTALLED_APPS = (
'django.contrib.auth',
'django.contrib.contenttypes',
'django.contrib.sessions',
'django.contrib.sites',
'django.contrib.admin',
'mysite.blog'
 )
```

Expert Included.

Falko provides expert, dedicated technical support for one of the most comprehensive server and storage product offerings in the industry.

He appreciates the Rackform nServ K501 because he knows AMD's Direct Connect Architecture ensures that all processor cores will work together with maximum efficiency and optimized memory performance now. And its 2 Dual-Core AMD Opteron™ 2000 Series processors are engineered for seamless upgradeability to Quad-Core later.

Falko knows that the evolution of technology is constant. He is impressed that the Rackform nServ K501's 24 hot-swap + 2 internal SATA drive system with redundant power supply delivers reliability without sacrificing capacity.

When you partner with Silicon Mechanics, you get more than a long-term investment in AMD technology—you get an expert like Falko.

ATTENDING LINUX WORLD?
Visit Silicon Mechanics in Booth 414 for the chance to win one of three valuable raffle items.

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

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

Notice the clear namespace distinction between my application (mysite.blog) and the applications that are included with Django (django.contrib.*).

Before we turn our Python code into SQL, we first should check to make sure it passes some basic sanity and validation checks. To do that, we go to our site's home directory, and type:

```
python manage.py validate
```

If all goes well, Django will report that there aren't any errors. Now that our model has been validated, we can use it to create SQL. The easiest way to do this is with the sqlall command to manage.py:

```
python manage.py sqlall blog
```

This produces the SQL output for our database driver (PostgreSQL, in this case). For example, this is the output that I see on my system:

```
BEGIN;
CREATE TABLE "blog_posting" (
    "id" serial NOT NULL PRIMARY KEY,
    "title" varchar(30) NOT NULL,
    "body" text NOT NULL,
    "publication_date" timestamp with time zone NOT NULL
);
COMMIT;
```

To their credit, the Django developers wrap the CREATE TABLE statement between BEGIN and COMMIT, ensuring that the table creation will take place in a transaction and will be rolled back if there is a problem. This isn't an issue when creating only one table, but if we have several models, it's always best to leave the database in a consistent state.

One way for us to use the output from sqlall to create tables is to copy it from the terminal window and then paste it, either into a file or into the psql client program. But, Django provides the syncdb utility to do this for us:

```
python manage.py syncdb
```

The output from this reassures us that all is well:

```
Creating table blog_posting
Loading 'initial_data' fixtures...
No fixtures found.
```

And, sure enough, now we can see that our table has been added:

```
atf=# \d blog_posting
 id               | integer                     | not null
                  default nextval('blog_posting_id_seq'::regclass)
 title            | character varying(30)       | not null
 body             | text                        | not null
 publication_date |  timestamp with time zone | not null
Indexes:
"blog_posting_pkey" PRIMARY KEY, btree (id)
```

Voilà! We now have a model that we can access via Python methods, but that exists in our relational database.

Inserting Data

Now that our data model is in place, let's see how we can work with it. Given that our model is brand new, and that there is no data currently stored in it, let's begin by adding some data to it.

In last month's column, we saw how each URL request in Django results in the invocation of a method. Which method is invoked depends on the settings of urls.py, a site-wide configuration file that tells Django what application and method should be associated with what URL.

One way to add data to our blog database, and to get some practice working with the various components of Django, is to do so via a view and template. Normally, I would demonstrate how to do this with an HTML form, but for space reasons, I use a simpler (and more contrived) way, inserting dummy data into the database.

The first step is to add a new line to the definition of the urlpatterns variable, defined in urls.py:

```
('^blog/add_dummy_data', 'mysite.blog.add_dummy_data')
```

Now, we can go to the URL /blog/add_dummy_data, and Django will invoke the blog.add_dummy_data method. The beginning of this method is quite simple, namely:

```
def add_dummy_data(request):
```

Listing 1. models.py for Creating New Dummy Posts

```
from django.template import Context, loader
from django.http import HttpResponse

from blog.models import Posting
from datetime import *

def add_dummy_data(request):
    p = Posting(title='Dummy 1 headline',  body='This is my
first blog post', publication_date=(datetime.now() -
timedelta(0, 0, 0, 0,1)))
    p.save()

    p = Posting(title='Dummy 2 headline',  body='This is my
second blog post', publication_date=datetime.now())
    p.save()

    return HttpResponse("Created blog posts.")
```

EmperorLinux

...where Linux & laptops converge

Portable

Since 1999, EmperorLinux has provided pre-installed Linux laptops to universities, corporations, government labs, and individual Linux enthusiasts. Our laptops range from full-featured ultra-portables to desktop replacements. All systems come with one year of Linux technical support by phone and e-mail, and full manufacturers' warranties apply.

Toucan T61/T61ws

ThinkPad T61/T61ws by Lenovo
- Up to 15.4" WSXGA+ w/ X@1680x1050
- NVidia Quadro NVS 140M graphics
- 1.8–2.4 GHz Core 2 Duo
- 512 MB–4 GB RAM
- 80–160 GB hard drive
- CDRW/DVD or DVD±RW
- 5.2–6.0 pounds
- 10/100/1000 Mbps ethernet
- 802.11a/b/g (54Mbps) WiFi
- Starts at $1650

Powerful

EmperorLinux specializes in the installation of Linux on a wide range of the finest laptops made by IBM, Lenovo, Dell, Sony, and Panasonic. We customize your choice of Linux distribution to your laptop and provide support for: ethernet, wireless, X-server, ACPI power management, USB, EVDO, PCMCIA, FireWire, CD/DVD/CDRW, sound, and more.

Rhino D830/M90

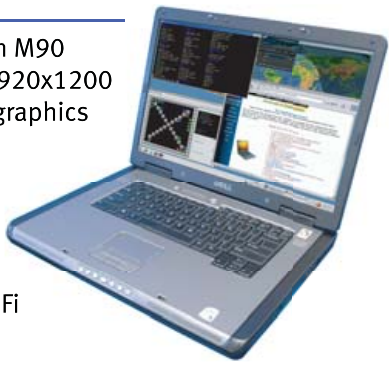

Dell Latitude D830/Precision M90
- Up to 17" WUXGA w/ X@1920x1200
- NVidia Quadro FX 3500M graphics
- 1.8–2.4 GHz Core 2 Duo
- 512 MB–4 GB RAM
- 60–160 GB hard drive
- DVD±RW or Blu-ray
- 6.3–8.6 pounds
- 802.11a/b/g (54Mbps) WiFi
- ExpressCard/EVDO
- Starts at $1445

Unique

EmperorLinux offers Linux laptops with unique features. Ruggedized Panasonic laptops are designed for harsh environments: drops, vibrations, sand, rain, and other extremes. ThinkPad tablet PCs are like other laptops, with an LCD digitizer for pen-based input both as a mouse and with pressure sensitivity for writing and drawing on-screen.

Raven X61 Tablet

ThinkPad X61 Tablet by Lenovo
- 12.1" SXGA+ w/ X@1400x1050
- 1.6 GHz Core 2 Duo
- 1–4 GB RAM
- 80–120 GB hard drive
- 3.8 pounds
- Pen/stylus input to screen
- Dynamic screen rotation
- Handwriting recognition
- X60s laptops available
- Starts at $2300

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

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

The name of the method is obvious from the configuration file. The number of parameters is determined by the number of parenthesized groups in urlpatterns.

Now what do we do? If we were dealing with raw SQL, I would suggest the following:

```
INSERT INTO Posting
    (title, body, posted_at)
VALUES
    ('Dummy 1 headline', 'This is my first blog post',
        NOW - interval '1 hour');

INSERT INTO Posting
    (title, body, posted_at)
VALUES
    ('Dummy 2 headline', 'This is my second blog post',
        NOW());
```

These will insert two rows into the Posting file: the first with a timestamp from one hour ago and the second with a current timestamp.

But, we don't want to use SQL. We want to use Python, creating objects that automatically map to these INSERT statements. So, it makes sense that all we have to do is create new instances of the Posting object, passing it appropriate parameters. And, sure enough, we can do that:

```
p = Posting(title='Dummy 1 headline', body='This is my
first blog post', posted_at=(datetime.now()
- timedelta(0, 0, 0, 0, 1)))
p.save()

p = Posting(title='Dummy 2 headline', body='This is my
second blog post', posted_at=datetime.now())
p.save()
```

Listing 2. views.py, with an Index Method

```
from django.template import Context, loader
from django.http import HttpResponse

from blog.models import Posting
from datetime import *

def index(request):
    postings = Posting.objects.all().order_by("-publication_date")
    output = ""

    for posting in postings:
        output += "<h1>%s</h1>\n" % posting.title
        output += "<h2>%s</h2>\n" % posting.publication_date.isoformat()
        output += "<p>%s</p>\n\n\n" % posting.body

    return HttpResponse(output)
```

If you are an experienced Python programmer, the above code shouldn't be very surprising at all. We simply are creating two new instances of Posting, passing arguments that will set the object's attributes. Then, we invoke the save() method on each posting, which presumably saves the posting to disk.

Finally, we finish our method with:

```
return HttpResponse("Created blog posts.")
```

With the method (shown in Listing 1) defined, start up the server:

```
python manage.py runserver  69.55.232.87:8000
```

Then, point the Web browser to the URL defined in urls.py, and get the message:

```
Created blog posts.
```

Next, check the database, just to be sure:

```
atf=# \x
Expanded display is on.

atf=# select * from blog_posting;

-[ RECORD 1 ]----+------------------------------
id               | 1
title            | Dummy 1 headline
body             | This is my first blog post
publication_date | 2007-06-15 16:13:34.609396-05

-[ RECORD 2 ]----+------------------------------
id               | 2
title            | Dummy 2 headline
body             | This is my second blog post
publication_date | 2007-06-15 16:14:34.675235-05
```

As you can see, we were able to create these new objects successfully and store them in the database.

Retrieving Data

Now that we've created these objects, let's see if we can retrieve and display them—a pretty typical thing to do if you write applications in Django. Because the most common thing you might want to do with a blog is display all of the postings in reverse chronological order, we write our index method to do that. If you still don't have an entry in urls.py for index, make sure there is a line that looks like the following in the definition of urlpatterns:

```
(r'^blog/$', 'mysite.blog.views.index'),
```

Now, we open up views.py to create our index method. The first task in that method is to get all the

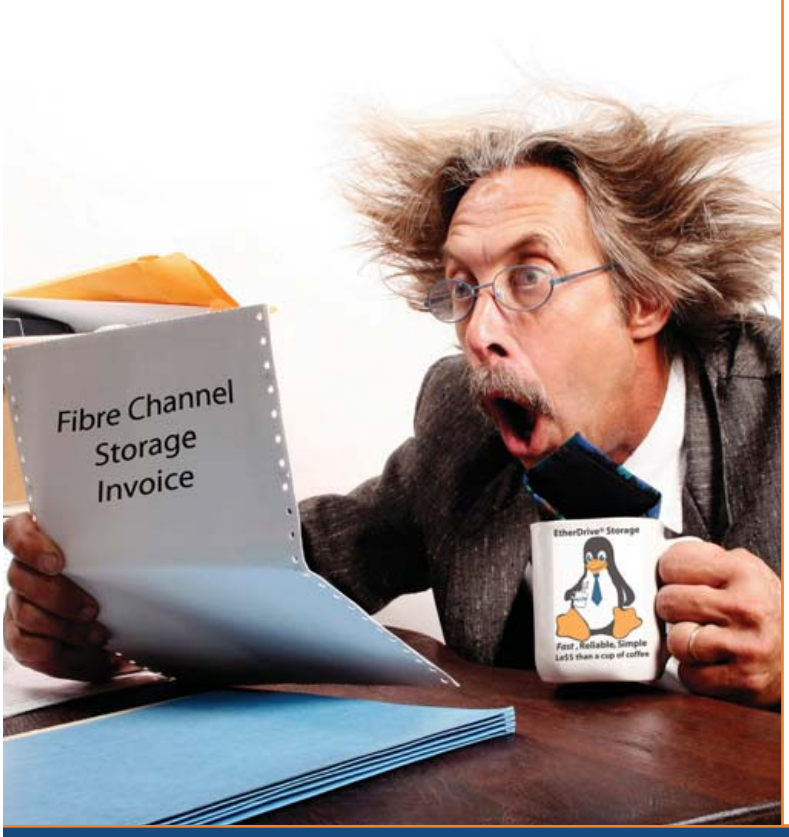

Are you Shocked

by the high cost of iSCSI & Fibre Channel SAN storage?

AoE is the answer!

ATA-over-Ethernet = ***Fast, Reliable, Simple*** storage.

www.coraid.com

EtherDrive® SRxxxx

- Fast & Flexible RAID appliances with slots for hot swap SATA disks
- Check out our full line of EtherDrive® Storage and VirtualStorage Appliances and NAS Gateways

Product Excellence Award Winner LINUXWORLD CONFERENCE & EXPO Best Storage Solution

1. Fast 10 Gigabit Ethernet Storage without the TCP/IP overhead!
2. Unlimited expandability, at the lowest possible price point!!
3. You want more storage…you just buy more disks – it's that simple!!!

Visit us at www.coraid.com

CORAID

The Linux Storage People

1.706.548.7200

www.coraid.com

postings. Django makes that trivially easy to do:

```
postings = Posting.objects.all()
```

This retrieves all the instances of Posting (which happen to be stored as rows in our database) and assigns them to the variable postings. This variable isn't a list, but an instance of a QuerySet object. We most likely will want to iterate over the QuerySet, but we can perform other operations on it, such as reordering it or retrieving selected elements.

We also can select particular items from the database. This is done with two methods: one called filter (which returns objects that match a restrictive function) and one called except (which does the opposite, returning objects that are false for a function). Both filter and except take a large number of parameters, built up dynamically by joining column names with various functions. The column name and function name are joined with a double underscore (__).

If you have worked with databases only from PHP or CGI programs, you will be surprised and impressed by the degree of automation Django provides.

For example, we can get only those postings from this year:

```
this_year_postings = Posting.objects.filter(
    publication_date__gte=datetime(2007, 1, 1))
```

Sure enough, this returns both of our postings. Because filter and except return QuerySet objects, we can chain them together, creating just the query we want in Python code.

But, what if we want only the most recent posting? If you're thinking there will be a "limit" feature, you've been working at the SQL level (or in Rails) for too long. Because QuerySets use lazy evaluation, you simply can say:

```
this_year_postings = Posting.objects.filter(
    publication_date__gte=datetime(2007, 1, 1))[0]
```

We similarly can order our objects by using the order_by method on them, which can be chained along with filter and exclude:

```
latest_posting = Posting.objects.filter(
    publication_date__gte=datetime(2007, 1,
1)).order_by('-publication_date')[0]
```

Notice that we put a minus sign (-) before the word publication_date. This tells Django we want to order the results in reverse.

Django has a wealth of such methods, giving both a great deal of flexibility in constructing your queries and a rich Python API that allows you to ignore the low-level SQL calls almost entirely.

Finally, we can get information out of our object as we would retrieve it from any Python object:

```
output += "<h1>%s</h1>\n" % posting.title
output += "<h2>%s</h2>\n" % posting.publication_date.isoformat()
output += "<p>%s</p>\n\n\n" % posting.body
```

If we put this all together, as shown in Listing 2, we'll have a view method (albeit without a proper Django template) that shows each of the blog postings.

You can try all of these database queries for yourself using Django's shell:

```
python manage.py shell
```

Using the Django shell, as opposed to the straight interactive Python interface, ensures that Django-related classes and paths are preloaded, making it possible to query and modify the database from within Python interactively. This is a good way to experiment with new code that you are thinking of adding to a view method, without having to place it in a file.

Conclusion

Django provides a high-level interface for the definition of database models using Python, rather than SQL. This high-level API permeates the framework, making it possible to work exclusively in Python. Moreover, the API includes many convenience functions and data types that make it relatively natural to work in this way. Creating database-backed Web applications with Django is dramatically easier and better than with most frameworks I've used, although it is similar in style to Ruby on Rails. Whether you should use Django or Rails is a matter of personal taste and also depends on what others in your organization are using, but there's no doubt that if you're a Python Web/database hacker, Django is worth a very serious look.■

Reuven M. Lerner, a longtime Web/database consultant, is a PhD candidate in Learning Sciences at Northwestern University in Evanston, Illinois. He currently lives with his wife and three children in Skokie, Illinois. You can read his Weblog at altneuland.lerner.co.il.

Resources

Django Documentation: **www.djangoproject.com/documentation**

Django Model API: **www.djangoproject.com/documentation/model-api**

Django Database API: **www.djangoproject.com/documentation/db-api**

Delivering a Successful Appliance Requires a Flexible Partner

"They bent over backwards"

"I can't think of too many other suppliers in any field who would've done what they did for us."

- Dave Troup, President and CEO of Solinus

Flexibility is Everything to Us

Your software on a customized platform, configured to your specifications, with your branding and has to work perfectly. **Got it.**
Need five or five thousand? **No problem.**
Logistics, shipping and support? **Handled.**

We're flexible so you can focus on your customers.
Or getting more of them.

Visit us at www.mbx.com or call 1-800-560-1195 today

MBX systems

MARCEL GAGNÉ

Still Searching for the Ultimate Linux Distro?

Why does a person install one Linux, then another, and then yet another? Because a person can, of course! Such is the nature of choice, and Linux gives you a choice...and what a selection.

What distribution are you loading up today, François? MCNLive? Very nice, and compact too. When you get a chance, you should copy it to your USB key. That way, you always can carry a live Linux distribution with you. *Quoi?* You're not sure if this is *the* one? I see. Yesterday, you were running OpenSUSE 10.2, and the day before you installed Debian Etch in the morning and Kubuntu Feisty in the afternoon. Last week, you managed Fedora Core 7, CentOS 5, Mandriva Corporate Desktop 4.0, Slackware Linux 12 and a half-dozen others. Are you having trouble finding something you like? You like them all but you just can't choose, eh?

Well, *mon ami*, I hate to interrupt this voyage of discovery, but I need that wine list I sent you a couple of days ago. What do you mean, you don't have it? It was on the machine where you are installing the distros and you erased the disk? You know, François, there are better ways to try out all these distributions. For the moment, leave what you are doing and head down to the wine cellar. I can see our guests approaching the restaurant now, and they will be here any second. Bring back the George du Beouf Cuvée Saint Valentin from the East wing. There are three cases right next to that old suit of armor you brought in last year. *Vite!*

Welcome, everyone, to *Chez Marcel*, where fine wine meets great Linux and open-source software. Please, sit down and make yourselves comfortable. I've already sent my faithful waiter on a quest to retrieve the perfect wine for tonight's menu. While we wait, I'm going to introduce you to an impressive parade of Linux distributions, courtesy of system virtualization. Best of all, you can keep running your current system while you take these others out for a spin.

Ah, glad to have you back, François. Please, pour for our guests while I introduce the first item on tonight's menu.

Fabrice Bellard's QEMU is a free, open-source, machine emulator and virtualizer. The reason for this distinction is that QEMU can emulate different machine types and hardware, but the performance, although not bad, can be greatly improved upon. Virtualization is achieved by using a kernel module that executes code on your system processor rather than emulating the processor. Another reason this is an important distinction to make is that QEMU also can emulate different processor architectures. For instance, if you wanted to run a SPARC or a PowerPC machine on your Intel processor, you could. Fabrice's Web site has a nice table showing the various processors that can be emulated. There also are some prebuilt QEMU images of different operating systems ready for download—you even can get FREEDOS and Minix if you like. What we're going to do though, is install Linux, lots and lots of Linux distributions.

Most modern Linux distributions come with QEMU, but the latest source always is available from **fabrice.bellard.free.fr**. The virtualization kernel module, however, usually requires that you download it from the site and build it yourself. Although you don't need it specifically, the performance improvements are dramatic and well worth the effort. Because this is a kernel module, you load it with the `modprobe kqemu` command.

Let's take a look at how QEMU works. For this first demonstration, I'm going to install Puppy Linux from an ISO image downloaded from the Puppy Linux Web site. Because Puppy is a cute little distribution with minimal space requirements, I create a relatively small disk image (a virtual hard disk) for it to live in. This is done with the `qemu-img` command:

```
qemu-img create puppy216.img 256M
```

The above command creates a raw format disk image by default. There are a few different image formats, most notably qcow2, which is a more portable image format, useful if you want to install that other OS—you know, the one from Redmond. Our next step is to install Linux into this disk image, which I do using this command:

```
qemu -cdrom ../isos/puppy-2.16.1-seamonkey-fulldrivers.iso \
   -hda ./puppy216.img -m 256 -boot d
```

Several interesting things are happening here, and I'll describe each one briefly. For starters, the -cdrom parameter is, in fact, the path to the CD-ROM image from which you are installing your distribution. If you were using a

Continuous Data Protection
The Future of Data Centers

Can your backup software do this?	R1Soft CDP Server	Acronis® True Image	EMC Retrospect ®
Daily Backups	✓	✓	✓
Hourly Backups	✓	Not Supported	Not Supported
Open File Backups	✓	✓	Not Supported
Bare-Metal Restore	✓	✓	Not Supported
Continuous Data Protection	✓	Not Supported	Not Supported
Restore Linux LVM	✓	Not Supported	Not Supported
Restore Linux Software RAID	✓	Not Supported	Not Supported
Easy To Use Web Interface	✓	Not Supported	Not Supported
Manage Thousands of Servers	✓	Not Supported	Not Supported
Control Panel Integration	✓	Not Supported	Not Supported
	$80 - $100 /server	$699 /server	You Can't Afford It

Data Centers serious about uptime and performance use R1Soft.

For more information visit: **www.r1soft.com** or call us at **800-956-6198**

R1Soft

Continuous Data Protection

For **LINUX & WINDOWS**

Copyright 2007 Righteous Software Inc All Rights Reserved.
R1Soft is a trademark of Righteous Software Inc. Other names may be trademarks of their respective owners.

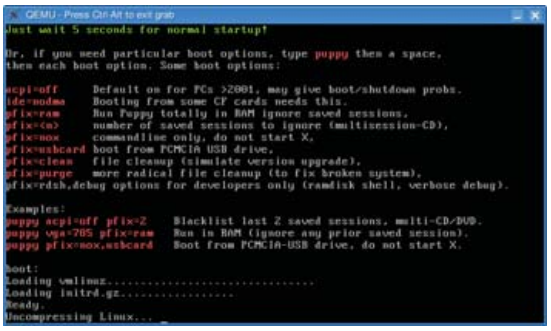

Figure 1. QEMU boots a Puppy!

physical CD-ROM, that path would likely be something like /dev/cdrom. The next parameter, -hda, defines the path to the disk image we just created. This is followed by the -m switch, which, in this case, allocates 256MB of RAM to the running session. Finally, we have -boot, which identifies the boot drive (our CD image), the so-called D drive. As soon as I press Enter, the Puppy Linux install starts (Figure 1).

From here, you simply follow the steps for a Puppy Linux installation (or those of whatever distribution you are installing). Before I show the result, I'd like to share one more parameter with you—the -std-vga parameter. This tells QEMU to simulate a standard VGA card with Bochs extensions. You also can use the -localtime parameter to tell QEMU to run your session using the local machine time, as opposed to the default UTC. Remember the kernel acceleration? Try using -kernel-kqemu. This is only a small sample. There are network parameters, USB parameters, file-sharing parameters (SMB), devices parameters and more. I'm covering only the basics here, so make sure you check out the man page for the qemu command. Remember also that my settings for this example imply a very small footprint distribution. To run something slightly larger, such as MCNLive (a great live distribution based on Mandriva), I would create a larger disk image (see Figure 2 for a shot of MCNLive in action).

When you start working in a QEMU session for the first time, you may find yourself wondering why you can't mouse out of the session. This is the first and most important lesson you will learn about QEMU keyboarding. Once you click in a QEMU session, it takes control of your mouse. Pressing Ctrl-Alt releases control so you can work on your master system. To return to the QEMU session, simply click anywhere on the window. The second most important keyboard command is Ctrl-Alt-F, which switches to and from full-screen mode.

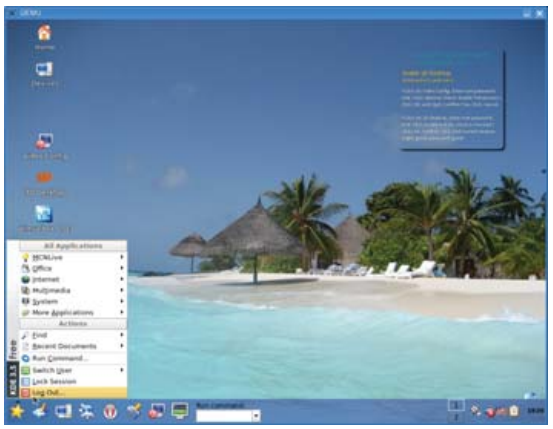

Figure 2. A Lovely Graphical Desktop, Courtesy of MCNLive and QEMU

Now, Puppy Linux and MCNLive are live distributions, but both offer an install icon on the desktop. Once a distribution is installed, starting the session is a little different, mostly because you won't be booting from the CD-ROM ISO image. In this example, I simply identify my hard disk image for QEMU and select the amount of RAM I want to use:

```
qemu -hda ./mcnlive.img -m 384
```

All this command-line work is easy enough, but there are dozens of qemu command parameters, not to mention tons of possible options or combinations of those. That's why we are going to look at some great graphical tools for working with QEMU, starting with Erik Meitner and Linas Zvirblis' Qemu Launcher.

Qemu Launcher provides a front end to basic as well as advanced features of the QEMU command suite. There are three tabs: Configurations, Launcher Settings and About. Most of the action takes place under the first tab. For starters, you can create a set of configurations based on existing QEMU images simply by entering a configuration name, a file location for your hard disk (the image created by qemu-img) and defining the memory you want to allocate. Click the Save button, and you are ready to start your virtual machine by clicking the Launch button on the lower right of the Qemu Launcher interface. In Figure

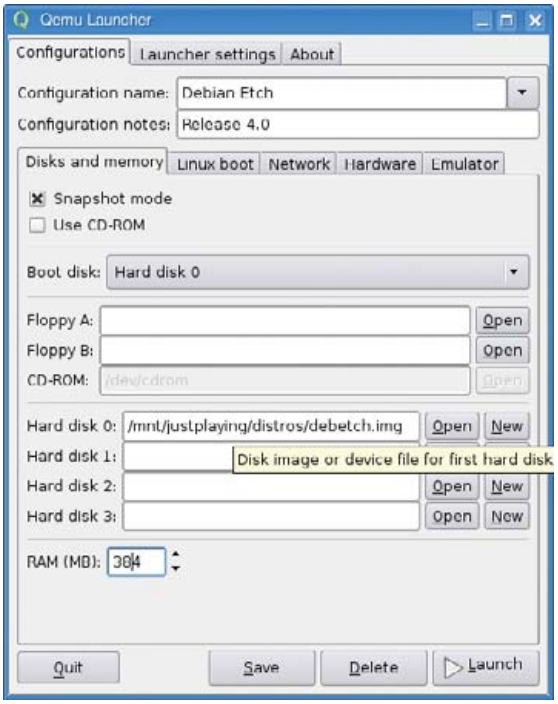

Figure 3. Adding an existing QEMU image under Qemu Launcher is easy.

intel Leap ahead™

MAXIMIZE PROCESSING PERFORMANCE AND MAXIMIZE RESPONSIVENESS.

COST-EFFECTIVE 2U SERVER

STARTING AT:
$1,599

- Intel® Server System SR2520SAXNA
- Quad-Core Intel® Xeon® Processor 5310
- 1024 MB FB-DIMM ECC Memory
- 400 GB 7200 RPM 16 MB Drive

1U CLUSTER SERVER

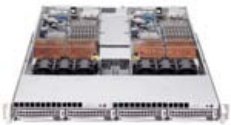

STARTING AT:
$1,999

- Two Nodes in 1U
- Dual-Core Intel® Xeon® Processor 5050
- 1024 MB FB-DIMM ECC Memory
- 250 GB 7200 RPM 16 MB Drive

1-800-576-7931
servers DIRECT
GO TO THE SOURCE

ServersDirect
1-800-576-7931
www.ServersDirect.com

ServersDirect can help you configure your next high-performance server system—contact us today!

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

intel Xeon® inside™
Quad-core. Unmatched.

Intel, the Intel logo, Intel. Leap ahead., the Intel. Leap ahead. logo, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.
*Other names and brands may be claimed as the property of others.
© 2007 Intel Corporation. All rights reserved.

3, I've added a profile for Debian Etch 4.0 to launch with 384MB of RAM.

I entered the path to the Debian Etch image, but I also could have clicked the New button next to the Hard disk 0 label and created an image using the pop-up dialog. My next step would have been to click the Use CD-ROM check box, select CD-ROM from the Boot disk drop-down list, and continue normally from there. As you can see, the Configurations tab is also a multitabbed affair that allows you to access other QEMU settings. Under Linux boot, you can specify an alternate kernel image or pass kernel parameters. The Network tab offers control over your network card (the default is to use a pass-through to your main system) and also provides access to a Samba share—useful should you want to run Windows. From the Hardware tab, you can choose an alternate video device, set the clock, define your hardware or turn on system sounds to name a few options. Finally, the Emulator tab gives you control over the QEMU session itself.

Incidentally, a recent alternative to QEMU is KVM (Kernel Virtual Machine), which works by taking advantage of processor-based hardware virtualization technology. On the Intel platform, this is known as VT, which AMD processors call their hardware virtualization, Pacifica. This means you can take advantage of Linux KVM to run your virtual machines with substantially better performance than you can with straight software virtualization. The catch, as you might expect, is that not every machine out there supports hardware virtualization. To find out whether your system is a candidate, execute this little snippet of code from a terminal window or shell:

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

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

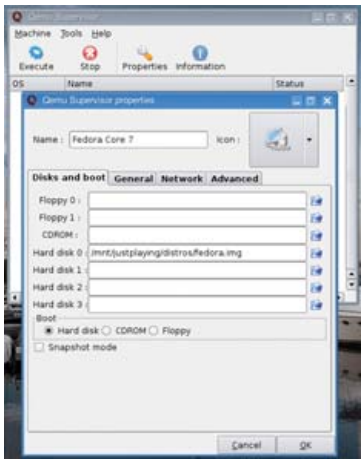

Figure 4. Now you can Ruby-up your QEMU sessions with Qemu Supervisor.

The next item on tonight's menu isn't as sophisticated or flexible (you can't create images), but those of you with a soft spot for Ruby may want to check out Mathieu Mast's Qemu Supervisor, a front end based on, you guessed it, the Ruby programming language. Aside from Ruby itself, you also need to have the Ruby GTK2 bindings, the libgtk2-ruby package and its dependencies. Once these prerequisites are met, it's only a matter of extracting the Qemu Supervisor source and running a `make install` from the source directory. To run the program, execute the command `qemu-supervisor`. From the interface that appears (Figure 4), you can choose to add QEMU images for launch.

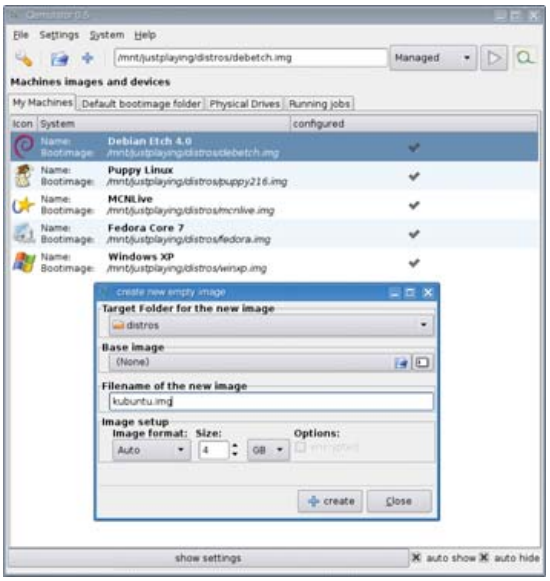

Figure 5. Qemulator makes QEMU easy with tools for image creation and more.

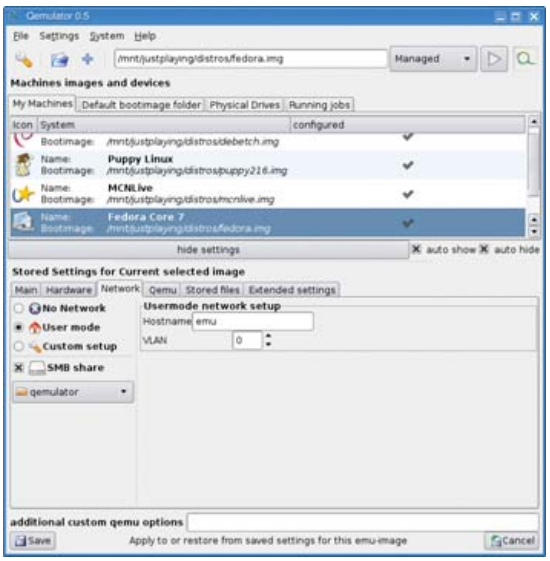

Figure 6. Easy access to all QEMU settings are a click away.

The final item on tonight's menu, Rainer Haage's Qemulator, is a must-see stop on your search for the ultimate QEMU front end to all your ultimate Linux (virtual) boxes. This impressive and feature-rich application covers pretty much every aspect of QEMU's operation, including creating disk images (Figure 5), installing distributions, monitoring running instances and so much more.

Like the other programs featured here, Qemulator provides easy access to existing QEMU images. Click the Plus sign at the top of the main window, enter the basics and save. To run a virtual machine, simply click the run arrow

at the top right. The main window lists existing images under the My Machines tab. Each of these machines can, in turn, be configured or changed by clicking the Show Settings button at the bottom of the main window (Figure 6). The tabbed view provides access to image settings, disks and other hardware, network settings and more. Several machines can run simultaneously (given adequate resources) with access to each under the Running Jobs tab.

There are some nice, somewhat hidden features as well. For instance, right-click on a machine in the My Machines list, and a small pop-up menu appears. From that menu, you can select Show command line—perfect for the curious who want to find out exactly what the qemu command is doing once all those settings have been tweaked. From that same menu, you can get information on the image itself, change settings and more.

To round things out, Qemulator even provides an Install Wizard to take you through the steps of loading up yet another distribution. Simply click System on the menu bar and select Install system.

Looking around the restaurant floor, I see that many of you are already discovering a downside of running all these virtual machines. Virtual or not, each instance of Linux (or whatever else) that you run, requires processor time and system memory. The more VMs you choose to run, the more resources you will be drawing on. I run a second distribution happily on my dual-core notebook, but a third is pretty much out of the question for this machine. The more resources you can provide, the faster your virtual machines will run and the more virtual machines you can run.

Ah, if only we could find some way to run multiple instances of our award-winning wine cellar. But then, François and I might not have the pleasure of sharing ours with you, *mes amis*. And now, sadly, the time has arrived, and we must head back to our non-virtual homes. However, we here at *Chez Marcel* would never dream of sending you off without a final glass of wine. François, please take a moment to refill our guests' glasses one last time. Raise your glasses, *mes amis*, and let us all drink to one another's health. *A votre santé! Bon appétit!*■

Marcel Gagné is an award-winning writer living in Waterloo, Ontario. He is the author of the all-new *Moving to Free Software*, his sixth book from Addison-Wesley. He also makes regular television appearances as Call for Help's Linux guy. Marcel is also a pilot, a past Top-40 disc jockey, writes science fiction and fantasy, and folds a mean Origami T-Rex. He can be reached via e-mail at mggagne@salmar.com. You can discover lots of other things (including great Wine links) from his Web site at www.marcelgagne.com.

Resources

QEMU: **fabrice.bellard.free.fr/qemu**

Qemu Launcher: **https://gna.org/projects/qemulaunch**

Qemulator: **qemulator.createweb.de**

Qemu Supervisor: **rubyforge.org/projects/qemu-supervisor**

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

The WFTL-LUG, Marcel's Online Linux User Group: **www.marcelgagne.com/wftllugform.html**

Unbeatable
Price, Quality, & Service

$1995

1U Supermicro 6015B-TV: (Customized System)
2x Xeon Quad-Core E5335 2GHz CPU,
4GB 667mhz FB-DIMM, 2x 250GB HD, Slim DVD...

$2395

2U Supermicro 6025B-TR+B: (Like above spec.)
2x Xeon E5335 CPU, 4GB RAM, 2x 250GB HD
w/ Redundant Power, 8 hotswap bays ....

*Free gadget w/ purchase!

Windows redhat SUSE fedora solaris RoHS COMPLIANT

Customizable system solutions since 1989

Tel: 1-800-875-8590 Fax:408-736-4151

KSC
KING STAR COMPUTER
1259 Reamwood Ave
Sunnyvale, CA 94089
www.kingstarusa.com
Email: sales@kingstarusa.com

intel Xeon inside™

Intel©, Intel© Xeon™, Intel Inside©, Intel© Itanium© and the Intel Inside© logo are trademarks or registered trademarks of Intel Corporation of its subsidiaries in the United States and other countries. Prices and availability subject to change without notice. Not responsible for typographical errors.

Baccarat Punto Banco, Part II

DAVE TAYLOR

Remember that *Baccarat* hands can't go to 11 or even 10 for that matter.

Last month, we visited the rarefied climes of the expensive private player rooms at Monte Carlo to learn about the elite gambling game *Baccarat*. The game is straightforward, but it's governed by a fairly complicated set of rules.

At its most rudimentary though, *Baccarat* has two players: a Banker and a Player, and each is dealt two cards, the rank of which are summed up to calculate the winning hand.

Any hand worth 10 or more is divided by 10, so a 9 + 6 isn't 15, but 5, and a hand of 3 + 4 = 7 would indeed beat it.

It gets complicated in the case of whether the Banker should draw a third card. The Player can draw one card only on any two-card point value of less than six, but the Banker can draw based only on the Player's initial hand, the card the Player drew and the Banker's hand value. It works like this:

1. If the Player didn't draw a card, the Banker draws a card on 0–5 and stands otherwise.

2. If the Player drew a 2 or 3 and the Banker has a total of 0–4, the Banker draws a card.

3. If the Player drew a 4 or 5 and the Banker has a total of 0–5, the Banker draws a card.

4. If the Player drew a 6 or 7 and the Banker has a total of 0–6, the Banker draws a card.

5. If the Player drew an 8 and the Banker has a total of 0–2, the Banker draws a card.

6. If the Player drew a 9, 10, face card or ace and the Banker has a total of 0–3, the Banker draws a card.

In all situations not detailed above, the Banker cannot draw a card, and the *Baccarat* hand will end either with the Player winning, the Banker wining or an *égalité*, or tie.

As you can see, it's a set of rules that apparently could be loved only by a computer programmer, so it looks perfect for us to expand the baccarat.sh game in this direction.

On to the Script

Last month's *Baccarat* game took everything into account up to, but not including, the Banker picking a third card, with all the complicated rules detailed above.

Running it as is, here's a typical sequence:

```
$ sh baccarat.sh

Welcome to Baccarat. You can choose to either
stake the player or dealer to win this game.
Remember, face cards are worthless and all point values
are modulo 10, with '9' the best possible hand value.
We're using Punto Banco rules with this simulation,
if you're already an expert...

** Player was dealt: 3 of Hearts, 10 of Clubs (hand value = 3)
** Dealer was dealt: 4 of Clubs, 10 of Spades (hand value = 4)

Player takes a card: 6 of Hearts (hand value = 12)

** Banker play rules yet to come...game ends inconclusively.
```

As you can see, instead of capturing the rules, I just have a "coming soon" stub. Think of it as, um, early alpha so far.

In fact, it's sufficiently in alpha that there's a bug displayed in the above sequence. Can you see it? How can the Player have a hand value of 12?

Fortunately, it's straightforward to see the problem:

```
playerhandvalue=$(( $playerhandvalue + $handvalue ))
```

This is just plain wrong. Because the line above hands all three Player cards to the handValue function, all this needs to be is the simpler:

```
playerhandvalue=$handvalue
```

Phew. Easily fixed. Now, where were we?

One thing that we need to ascertain easily is whether the Player took a card. This can be done by actually testing the number of cards in the Player hand, but instead I simply create a new variable, playerDrewCard, setting it to zero on initial deal and flipping it to one if the Player takes another card.

Now, the first rule can be captured like this:

```
if [ $playerDrewCard -eq 0 ] ; then
  if [ $bankerhandvalue -lt 6 ] ; then
    # Banker draws a card
```

The Straight Talk People℠
SINCE 1991

ABERDEEN

BLADE KILLER
2 SERVERS IN 1U

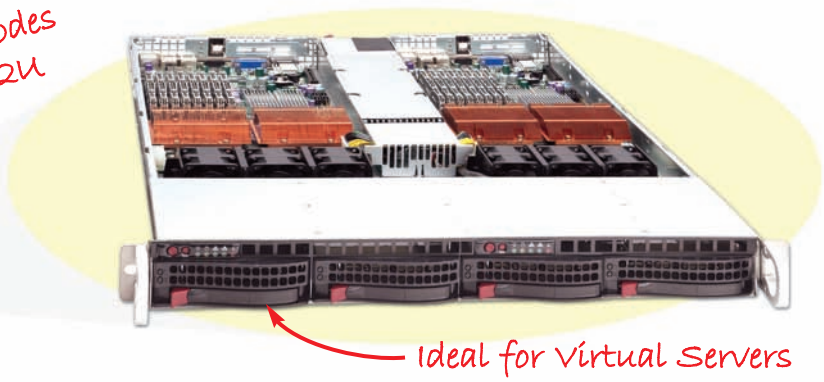

Features and benefits:

- **Higher Density than Blades** (Yes, really) — Two dual quad-core servers in a single 1U. Up to 84 nodes / 672 processor cores in a 42U rack.
- **Lower Power** — Superior power utilization. Increased power supply efficiency.
- **Redundancy** — High redundancy clusters. Independent servers instead of single points of failure found in blade solutions.
- **Standardization** — Industry standard architecture. No proprietary blade backplane/architecture.
- **Cooling** — Individually cooled 1U chassis vs. several blades in a single large chassis.
- **Complete Customization** — Every node is considerably more customizable than any blade.
- **Hot Swappable** — Swap an entire dual server without interrupting network flow or data array access.
- **Lower Cost** — Substantially lower cost than comparable blade solutions.

	IBM BLADECENTER® H	ABERDEEN STIRLING 122
	9U Blade Solution with 14 Blades	**7U Server Solution with 14 Nodes**
Description	9U IBM BladeCenter H chassis with 14 IBM BladeCenter HS21 blade servers	7 1U Aberdeen Stirling 122 Servers, each with two server nodes
Processors per blade / node	Two Quad-Core Intel® Xeon® processors E5310 1.60 GHz with 8 MB cache and 1066 MHz FSB	Two Quad-Core Intel® Xeon® processors E5310 1.60 GHz with 8 MB cache and 1066 MHz FSB
Memory per blade / node	2 GB ECC DDR2 667 MHz memory (max 32 GB / 4 DIMM)	2 GB ECC DDR2 667 MHz memory (max 32 GB / 8 DIMM)
Hard drives per blade / node	Two 36 GB 10,000 rpm internal 2.5" SAS drives (max 146 GB without optional SIO blade)	Two 36.7 GB 10,000 rpm hot-swap 3.5" SATA drives (max 1.5 TB)
Maximum blades / processors per 42U rack	56 blades / 112 processors (448 cores) in 4 x 9U chassis (6U empty)	84 server nodes / 168 processors (672 cores) in 42 x 1U chassis (0U empty)
Warranty	3 year on-site limited warranty for parts and labor	5 year limited warranty for parts and labor
Total price for 14 blade / node solution	$70,161	$36,495
Price per blade / node	**$5,012**	**$2,607**

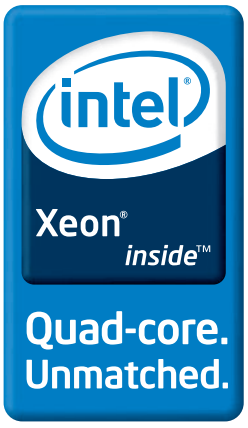

Intel, Intel Logo, Intel Inside, Intel Inside Logo, Pentium, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Other trademarks are of their respective owners. Prices and specifications for IBM products obtained from www.ibm.com on Januray 30, 2007. For terms and conditions, please see www.aberdeeninc.com/abpoly/abterms.htm. Ij018

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

```
    banker[$nextbankercard]=${newdeck[$nextcard]}

    handValue ${banker[1]} ${banker[2]}
      ${banker[3]}
    bankerhandvalue=$handvalue

    echo -n "Banker takes a card: "
    showCard ${banker[$nextbankercard]}
    echo "$cardname (hand value =
      $bankerhandvalue)"
  fi
else
```

Notice here that if the Player drew a card and the Banker's hand is worth six or more, there is no else clause, and correctly, there is no additional action before the end-game winner is determined.

The else at the bottom, however, is our gateway to the situation where neither the Banker nor Player has an 8 or 9 (which is already captured in the game) and where the Player took another card. Now, it's just a set of conditionals.

As you can see, it's a set of rules that apparently could be loved only by a computer programmer, so it looks perfect for us to expand the baccarat.sh game in this direction.

To make things easy though, let's have the Player's drawn card's rank handy:

```
pdcRank=$(( ${player[3]} % 13 ))
```

Unfortunately, although the man page for test indicates that you can group complicated tests logically with parentheses, pragmatic reality demonstrates that it's far less portable than we might desire, so instead of a nice (A or B) and C statement, we'll break each rule into two if statements, like this:

```
if [ $pdcRank -eq 2 -o $pdcRank -eq 3 ] ; then
  if [ $bankerhandvalue -lt 5 ]; then
```

In this case, we're testing the card the Player took (card #3) and testing that the Banker's hand is below a certain value. If both are true, we have condition #2, above.

Now, it's just a matter of having a sequence of these tests in a row to build all the rules necessary for *Punto Banco Baccarat*.

In all its crazy-nested-if-statement glory:

```
if [ $pdcRank -eq 2 -o $pdcRank -eq 3 ] ; then
    if [ $bankerhandvalue -lt 5 ]; then
      dealBanker;
    fi
  elif [ $pdcRank -eq 4 -o $pdcRank -eq 5 ] ; then
    if [ $bankerhandvalue -lt 6 ] ; then
      dealBanker;
    fi
  elif [ $pdcRank -eq 6 -o $pdcRank -eq 7 ] ; then
    if [ $bankerhandvalue -lt 7 ] ; then
      dealBanker;
    fi
  elif [ $pdcRank -eq 8 -a $bankerhandvalue -lt 3 ]
  then
    dealBanker;
  elif [ $pdcRank -eq 9 -o $pdcRank -eq 0 ] ; then
    if [ $bankerhandvalue -lt 4 ] ; then
      dealBanker;
    fi
  fi
```

Finally, we can play *Baccarat* on our handy Linux box:

```
> sh baccarat.sh

Welcome to Baccarat. You can choose to either
stake the player or banker to win this game.
Remember, face cards are worthless and all
point values are modulo 10, with '9' the best
possible hand value. We're using PuntoBanco
rules with this simulation, if you're already
an expert...

** Player was dealt: 6 of Clubs, 8 of Hearts
(hand value = 4)
** Banker was dealt: King of Spades, Ace of Hearts
(hand value = 1)

Player takes a card: 9 of Hearts (hand value = 3)
Banker takes a card: 5 of Clubs (hand value = 6)

Play is complete. Banker wins
```

Of course, now you can run this a few million times and calculate the odds of the Banker winning versus the Player winning versus the tie situation and be a cool, calculated gambler next time you're in the south of France. Helpful, eh?■

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

The baccarat.sh script is available on the *Linux Journal* FTP site: **ftp.linuxjournal.com/pub/lj/listings/issue161/9780.tgz**.

RouterBOARD™ 100 Series

made by routerboard.com

RouterBOARD 133

Low cost AP/CPE $89
RouterOS L4 Firewall/Routing
MIPS32 4Kc 175MHz embedded
32MB SDRAM
64MB onboard NAND storage
Three 10/100 ethernet ports MDI/X
Three miniPCI Type IIIA/IIIB slots
Power: 9-28V DC; Overvoltage protection
PoE: 12-28V DC (no power over datalines)
117mm x 105mm (4.61 in x 4.13 in), 115g

RouterBOARD 133C

Low cost CPE $59
RouterOS L3 Firewall/Routing
MIPS32 4Kc 175MHz embedded
16MB SDRAM
64MB onboard NAND storage
One 10/100 ethernet ports MDI/X
One miniPCI Type IIIA/IIIB slots
Power: 9-28V DC; Overvoltage protection
PoE: 12-28V DC (no power over datalines)
117mm x 105mm (4.61 in x 4.13 in), 79g

JON "MADDOG" HALL

Education

Lessons on evangelizing free and open-source software over a round of libations.

"How can I convince my professors to use free and open-source software (FOSS) in their classes?" asked Dennis, a former Pollywog that is now going to a university in Brazil. Dennis is old enough now to join me legally for libations at the *Alideia dos Piratas*, the beach-side restaurant and bar where I hang out, although he often hints that he had been "libating" much longer than the legal age limit allows.

"The first thing you need to do", I answered, "is create compelling arguments for the use of FOSS in the classroom. These arguments might include freedom from proprietary licensing and licensing's inherent costs and issues, the fact that the FOSS source code allows you both to use the software as a tool and to see how the software arrived at its answer. Publishing research results also is easier with FOSS, because you can freely publish source code for your work, not just talk about the work conceptually. FOSS allows schools to have better control over the software used on their systems, especially when they need to upgrade software and hardware. You should try to stick to practical reasons that are easy to defend, including the fact that FOSS is used in many commercial environments, which means that students will use it when they leave school.

"Present these arguments to the faculty and administration in a calm, professional way. If they do not respond favorably, ask their reasons for using proprietary software and ask whether you can address those reasons in a follow-up meeting.

"Second, make it as easy as possible for them to use FOSS. Remember that professors are like a lot of other people, and they are very busy just trying to keep up with all the other issues in their jobs and lives. Developing new courses and course material takes time and research, and although you can look at it as part of their jobs, you also can be assured that they are busy with other aspects of their jobs.

"One tactic would be for you to take an existing course, such as compiler design or database theory and transform it to using FOSS. You could find (or create) a bibliography of all the open-source projects that you can locate on the Web having to do with that subject. See if there are any books that relate to teaching that topic, which are oriented toward FOSS rather than closed-source products.

"In the case of the operating system, a wealth of projects using FOSS can be used to help teach a course in OS design—not only the Linux kernel, but the *BSD (OpenBSD, FreeBSD, NetBSD) kernels, FreeDOS, TinyOS and a variety of others. There also are several on-line talks and papers discussing aspects of FOSS OS design that can be useful, and for the more popular operating systems, specific books describe the kernel and the way it works.

"In the case of database design, a series of different projects also are useful—everything from MySQL and PostgreSQL to Berkeley DB (although an Oracle product, it is still open source), Firebird, flat-file databases and libraries for other types of data accesses, such as ISAM.

"Accumulating the bibliography for the course might be enough to gain the professor's interest. Many times professors have heard the terms free software or open source, but they have not had the time to investigate how much software and resources are out there with respect to their own courses.

"If the professor still doesn't accept using FOSS to teach the course, you could take the next step of actually devising a set of notes for the course as you study the supplied material. Nothing helps you learn the subject matter as much as having to try to organize it and explain it to others, and creating a set of class notes and presentation slides based on FOSS would help you learn the underlying information very completely."

Dennis looked at me with a sort of horror on his face. "That sounds like a lot of work", he said.

"It could be", I agreed. "Fortunately, you may not have to do all the work yourself, as Googling the Internet using the terms course notes and operating system turned up several editions of class notes that were under the GPL or copyable with attribution. You also might enlist some other students in the class to help by doing different chapters and then merge the work back into a single document.

"I took compiler theory twice, once as an undergraduate and once for my MS in Computer Science, yet I feel that I really did not know compiler theory until I had to teach it to others. Now, more than 20 years after giving the last formal class on that subject, I feel I easily can explain fundamentals of how a compiler works. Knowing difficult material that thoroughly is kind of scary in a way."

Both Dennis and I sat in silence for a couple of seconds contemplating this thought, then we both ordered another libation.

"Can other areas of the university also use FOSS methods to teach their courses?" asked Dennis.

"Of course. There are thousands of FOSS programs that students and faculty can use in engineering, the sciences and other aspects of education. In addition, there

are freely available pieces of information like Project Gutenberg, which houses more than 17,000 books and articles whose copyright have expired. I have a friend who teaches English as a second language but did not know of this treasure of on-line books that she could use for her students free of charge.

"Some disciplines, such as civil engineering, use a lot of maps and mapping software. The use of software from projects, such as FreeGIS (**www.freegis.org**) and Open Source GIS (**opensourcegis.org**), allow these academic fields to use and understand mapping software. Hospital administration and management courses can benefit from the OpenVista Project, a complete hospital administrative system used by the United States Department of Veterans Affairs that was then made freely distributable. Management courses at universities can benefit from sophisticated FOSS project and portfolio management systems, such as project.net (**www.project.net**), both to train students how to use such tools and show students how those tools work."

"What about the administration of our school, are there any FOSS programs that can help them?" asked Dennis.

"Yes. Schools often spend huge amounts of money for administrative software that barely works. I know of one major university that was going to spend ten million dollars to buy and implement an administrative system.

"Another university in Australia had spent 17 million Australian dollars on a commercial system, which (after five years) did not work at all. The Australian university was told that it would cost an additional five million Australian dollars and another three years of development to get a 'vanilla' system, one that was not tuned to their school, but at least worked. Maybe.

"Compare all of that to SAGU, a system that was created in Brazil as a FOSS project, which is now used by more than 60 universities throughout the world, having been translated from its native Portuguese to Spanish and English. Another interesting project is that of the Kuali Foundation, in its effort to create FOSS financial software for universities.

"Although not exactly an administrative program, the Sakai Project (**www.sakaiproject.org**) is a FOSS system that is being developed by a consortium of companies and educational institutions for learning and research collaboration. There's also Moodle, a well-known CMS system for use in education. All of these, and more, are available for use and can save huge amounts of money for budget-strapped schools and universities."

Dennis finished his drink, thought about what I had said, and asked the question, "What happens, if after all this work to try to convince my university to use FOSS, they still refuse to use FOSS?"

"The efforts you take to research and make the argument about using FOSS software would not be wasted. Perhaps you will learn more about your subjects by making these efforts, and that can never hurt. I know a young man who started a FOSS project to write an operating system kernel when he was going to college in Helsinki, Finland...." ■

Jon "maddog" Hall is the Executive Director of Linux International (www.li.org), a nonprofit association of end users who wish to support and promote the Linux operating system. During his career in commercial computing, which started in 1969, Mr Hall has been a programmer, systems designer, systems administrator, product manager, technical marketing manager and educator. He has worked for such companies as Western Electric Corporation, Aetna Life and Casualty, Bell Laboratories, Digital Equipment Corporation, VA Linux Systems and SGI. He is now an independent consultant in Free and Open Source Software (FOSS) Business and Technical issues.

ASA COMPUTERS
The Expert on Customized Servers !
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. "Runs both 32-and 64-bit applications simultaneously".

Your Custom Applicance Solution!!

"Let us know your Needs..."

"We will build you a Solution...."

AMD Opteron(TM) Value Server Starts at $847

- 1U 14" Deep 260W.
- AMD Opteron 140 CPU.
- 512MB PC 3200 DDR ECC Unbuffered.
- Support upto 8GB DDR RAM.
- 40GB SATA Hard Disk.
- 2x 10/1000 Mbps Lan.

Quad AMD Opteron(TM) Server Starts at $2,812

- 1U AMD Opteron Model 840.
- 2GB Memory. Max 128 GB
- Supports upto 64GB FBDIMM.
- 80 GB SATA II Hotswap Hard Drive.
- 2xIntegrated Dual 10/1000 LAN.

Dual AMD Opteron(TM) Storage Starts at $4,120

- 5U Dual AMD Opteron Model 246.
- iSCSI or NAS Software Options.
- Support upto 18TB of Storage.
- Fail Hard Drive LED Indicator.

Dual AMD Opteron(TM) Storage Starts at $8,445

- 8U AMD Opteron Model 246.
- 4TB of Storage (36TB Max).
- 1GB RAM
- 2 x 10/100/1000 Gigabit LAN.
- NAS or iSCSI Software Options.

Why Do Business With ASA?

"We Provide Approved EVAL Server..."

Since 1989, ASA has served customers like Cisco, Juniper, Caltech, Fermilab and most Universities.
We provide a total custom solution with OS of your choice.
Excellent pre and post-sales support.
"Reliable hardware at the most competitive prices".
Please call or contact us for your next hardware purchase.

AMD 64 Opteron

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

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

Navigating with the Nokia N800

DOC SEARLS

Putting Navicore and Maemo Mapper through the paces.

I'm a geo-freak. Geography and geology are two of my favorite subjects. Maps are among my favorite reading material. I can sit and study a map with great interest for an hour or more—even if it describes a place I'll never visit. I can fly from Los Angeles to Paris and stare outside at the ground for the whole distance, looking at the map that is the world itself, constantly hungry to know more about it. I consider GPS (Global Positioning System, **en.wikipedia.org/wiki/Global_Positioning_System**) one of the greatest inventions in the history of our species. I carry a handy little Garmin GPS unit (the excellent eTrex Vista CX, **https://buy.garmin.com/shop/shop.do?cID=145&pID=326**) with me nearly everywhere I go, and consider it no less essential than a cell phone.

So, you can guess my response when I got an e-mail from our Nokia contact, asking if we'd like to check out Navicore's GPS for the N800. I said yes and had a box in my hands the next day—which was good because we were already up against the deadline for this issue of *Linux Journal* (in which Jim Thompson and I give a big thumbs-up to the N800 as the Best Handheld Linux Device).

The N800 (and the 770, which preceded it) already has available GPS software in Maemo Mapper, free (GPL'd) software developed and continuously improved by John Costigan and a dedicated crew of programmers.

Navicore's and Maemo Mapper's GPS approaches are very different, but together they validate the N800 as a platform of potentially unlimited flexibility and appeal. The potential for Linux in an open mass-market portable computing device has been obvious for a long time, but reality has always lagged behind promise. But, now that we see companies like Navicore providing consumer-grade hardware and software kits for the N800, we know embedded Linux has passed an important milestone.

I've been using both the Navicore and Maemo Mapper GPS systems as much as possible on the N800 during the week I've pushed back the deadline. In the following order, here are the routes where I checked out both systems, alternating between the two:

- Drives around Santa Barbara and the nearby countryside.
- Flights from Santa Barbara to Los Angeles and Oakland.
- A bus from Oakland's airport to the nearest BART station.
- BART from the station to the tunnel under the Bay.
- Walks and taxi rides in San Francisco.
- A red-eye flight from Oakland to Houston.
- Taxi rides in Houston.

Unfortunately, I was unable to use it on the flight from Houston back to Los Angeles, or the drive from that airport to Santa Barbara, because the water bottle in my bag (bought after passing through security) leaked on the N800 and rendered it useless until it recovered the next day, presumably because it dried out. (The fact that it worked fine after being soaked amazes me.)

I'll start with a report on Navicore, which was the first of the two systems I put to use.

The Navicore package is a box of about the same size and shape as the N800's. Contents include:

- A Nokia battery-powered LD-3W GPS receiver (about the size of a Zippo lighter or a piece of hotel-room soap).
- AC power charger (US standard).

The Navicore, which stores its maps on a 2GB removable-SD card, was able to track progress along the ground during the flight from Santa Barbara to Los Angeles.

SUPERMICRO®

UIO

Universal I/O

Maximized I/O Expandability & Flexibility

- ✔ 3 / 7 PCI-E/X Cards for 1U/2U
- ✔ Highly Upgradeable
- ✔ High Efficiency Power (up to 90%+)
- ✔ Multiple Expansion Card Options (SAS RAID 5, 10-G, IB...up to 20 choices)
- ✔ Simplify Your Inventory Management

Gives you the freedom to build exactly what you need!

By supporting more add-on cards and I/O bandwidth, Supermicro 1U or 2U UIO systems meet the requirement for many applications that previously required a 2U or 3U system.

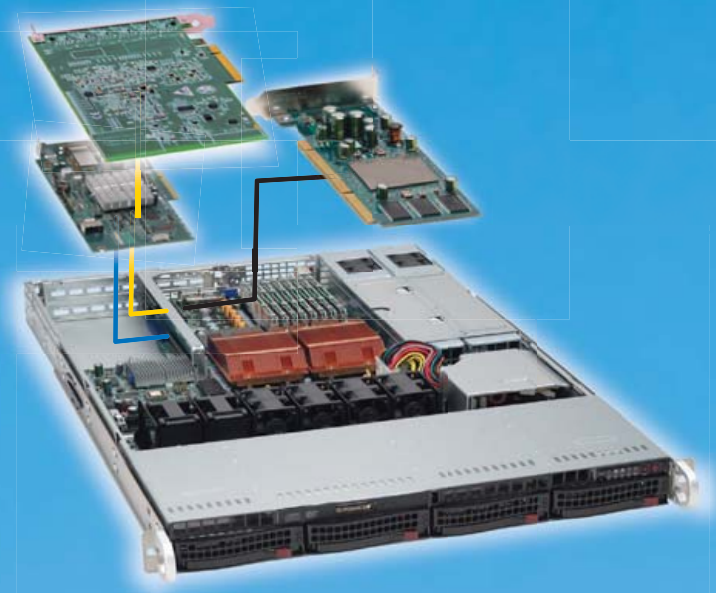

Optimized Chassis Solutions

SC815TQ-R650U

SC825TQ-R700U

UIO Motherboard Options

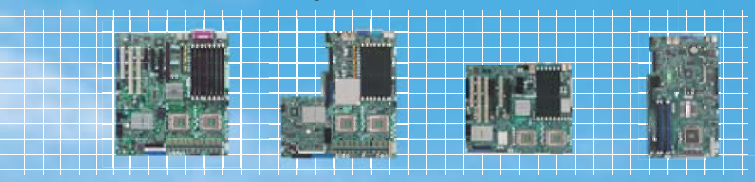

X7DBN | X7DBU | X7DAL-E+/X7DAL-E | PDSMU

Optimized SAS and RAID Cards

AOC-USAS-L4iR

AOC-USAS-H8iR

AOC-USAS-S4i

MX64MK00

TM

intel®

Xeon® inside™

For more information visit us at www.supermicro.com

* Our new generation power supply efficiency measures 90%+ under a typical loading operation.

AMAX	Arrow Electronics	ASI	Bell Micro	Ingram Micro	MA LABS	Synnex	Tech Data
1-800-800-6328	1-888-427-2250	1-800-2000-ASI	1-800-232-9920	1-800-456-8000	1-408-941-0808	1-800-756-5974	1-800-237-8931
www.amax.com	www.arrownacp.com	www.asipartner.com	www.bellmicro.com	www.ingrammicro.com	www.malabs.com	www.synnex.com	www.techdata.com

© 2007 Super Micro Computer, Inc. Specifications subject to change without notice. All other brands and names are the property of their respective owners.
Intel, the Intel logo, Intel inside, the Intel Inside logo, Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

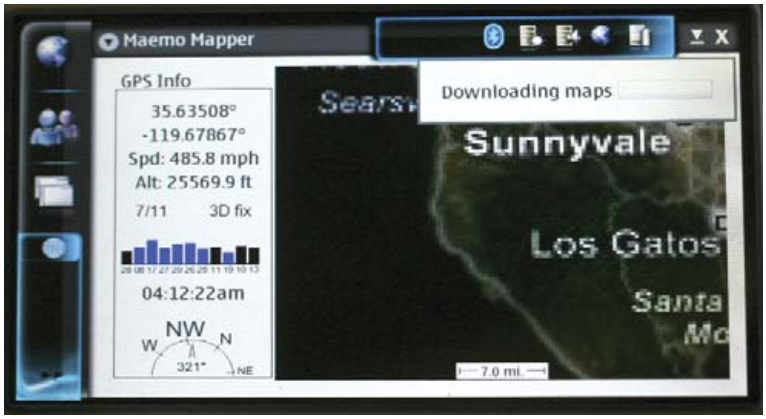

Although Maemo Mapper was able to use the GPS receiver and give latitude/longitude/altitude/heading readings, it wanted to download maps from the Web, which doesn't work on a plane without a Net connection. Without a fresh download, the map doesn't match the latitude/longitude readings on the left.

Trying Maemo Mapper on the LAX-OAK flight. The only useful data are the GPS readings. These include altitude, which the equivalent Navicore display does not. On the other hand, Navicore has its maps in memory and does not need to download them from the Net (which Maemo Mapper is failing to do in this case, because there is no Net connection).

Here we're in Navicore's 2-D view, which shows BART's railroad tracks as dotted lines. The 3-D view shows only roads (presumably because your car or truck is not on rails).

- A 2GB mini-SD card with maps covering the entire US and Canada.
- Car power adapter (for the lighter socket).
- Mounting hardware.
- A (Windows-only) DVD with the same data as the memory card.
- Support card with numbers for many countries.
- Minimal instructions in many languages.

The N800 has room for two SD cards—one internal (under the battery) and one external, revealed when the unit's stand is angled outward. I had already put a 1GB card inside the device, so I slid the 2GB Navicore card into the external slot. It was easy and intuitive to install the software using the N800's application manager.

I got the license activated, chose English, and specified my region as USA-West. Then, I put the GPS receiver by the window while it charged and pressed the button to turn it on.

Getting the N800 and the GPS to see each other via Bluetooth was easy and proceeded in the manner typical of Bluetooth devices (the passcode was the usual four zeros). Once the two devices were associated, activating the GPS through the Navicore UI was a one-button move. (And it continued to work fine for Maemo Mapper as well.)

Navicore's Navigation Kit is made to work primarily on roads as a car navigation system. You can find locations and then plan routes, which you also can clear. You can find detours as well, which can be handy if you run into traffic. You can choose to find locations by address, city, favorites, points of interest (POIs) or by latitude/longitude coordinates.

The screen centers on your location, which can be displayed in 2-D or 3-D (the latter being the preferable default in car navigation systems these days). Holding down the select button allows you to toggle between Browse and Navimode. Browse lets you move about the map, while Navimode gives you choices of views from your current location. Both make use of the four directional buttons surrounding the select button.

On the screen, a gear icon brings up a control panel with choices for the following:

- Guidance volume (four levels from mute to loud).
- Routing mode (car, bike, foot).
- Route type (shortest, fastest).
- Display mode (2-D, 3-D).
- Colour theme (daylight, night, brightlight).

Across the top are icons for satellite data and Bluetooth device readings. The satellite data include latitude/longitude, heading, speed, compass direction, satellite locations and signal strengths. It does not display elevation (one of my faves), presumably because the setup is meant for drivers rather than pilots. The device readings tell you the name and MAC address of the GPS receiver and its battery level, and there's a button for activating or deactivating the receiver.

Maps have icons for parking, hotel/motel, airport, amusement park, campground, theater, casino, college, concert hall, Wi-Fi access points and other points of interest. Clicking on one brings up its name and other possibly helpful information. (For example, it tells you whether the access point is a pay one.) If you click on an access point, your choices include adding it to favorites and navigating your way there (at the end of which the voice says, "Now get out of the car!").

Back on the hardware side, buttons for back, home and menu all do their jobs and work logically.

My car is a 2000 Volkswagen Passat station wagon. Although it's not a small car, its windshield leaves no room for the bulky N800 mounting hardware. I've used other GPS systems in cars before, especially the Garmins that

Navicore's zoomed-out view to the ground about midway between Los Angeles and Oakland. Note the speed reading in the upper left.

car-rental agencies offer. All are considerably smaller than the N800/mounting hardware combination. So, I'm sorry I was unable to test the mounting setup because its construction appears to be very good.

I did use the holder for the N800 alone, however,

PGI® CDK™
MPI Debugger/MPI Profiler

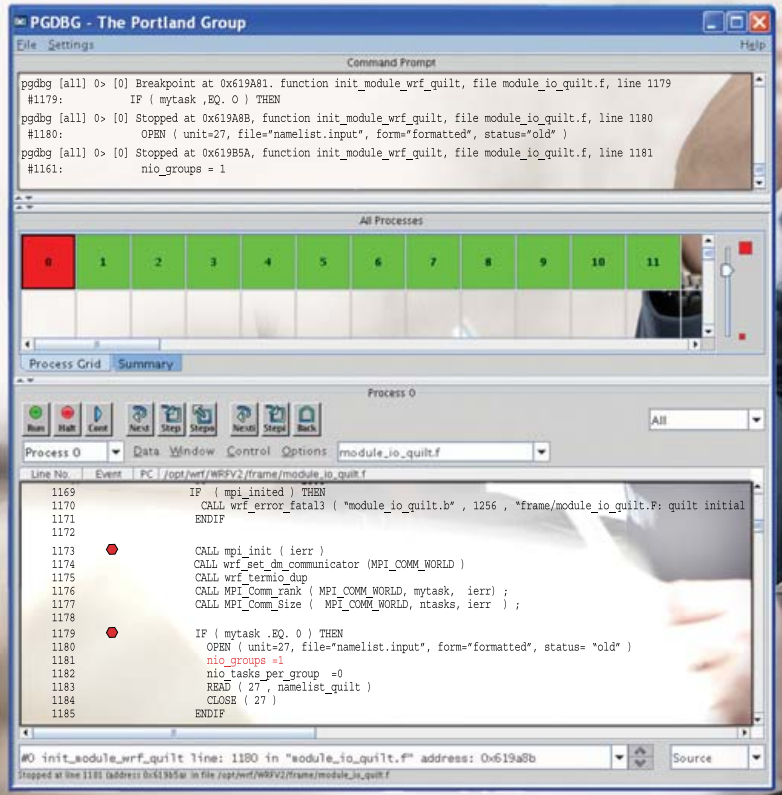

www.pgroup.com/cdk

The Portland Group, Inc. is an STMicroelectronics company. PGI and CDK are trademarks or registered trademarks of STMicroelectronics. Other brands and names are the property of their respective owners.

Maemo Mapper will create a track that leaves a red line. Here I'm taking a bus to a BART station from the airport in Oakland. The GPS receiver is in my pocket.

because it provided a nice protective frame around the device and enlarged it enough to make it fit nicely between the two front seats. The location of the GPS receiver didn't seem to matter. Whether it was on the seat, the dashboard or sandwiched between the car's roof shade and the dark glass moon-roof cover (a clever hack on my part, I thought), it worked fine.

We gave the Navicore our first test when we drove from our home in Santa Barbara toward Los Olivos, about 25 miles up Highway 154 in the Santa Ynez Valley (the wine country where most of the movie *Sideways* was filmed). Even when set for the shortest route, Navicore still wanted us to get on Highway 101 to reach 154, which made no sense from our start point, or any point along the most direct and obvious route. As we drove directly toward 154, the Navicore voice kept telling us to "Try to make a U-turn if possible". After we were on 154, however, it worked fine. Later, when we set it up to find our house on the return trip, it again wanted us to take the long "faster" way home, telling us to turn around, over and over again.

In San Francisco, it did a better job when set for walking. Here I put the GPS in my briefcase and held the N800 in my hand. I looked weird walking around that way, but the setup did a nice job of navigating me from one hotel or restaurant to another.

I also had fun using it on the airplane between Los Angeles and Oakland, and in the first half hour of a red-eye from Oakland to Houston. It was fun watching the plane "drive" across a close-up road map at speeds

Expert Included.

Forrest and the Product Development team never stop thinking about the problems facing today's IT administrators. Power efficiency, compute density, green computing, reliability, and serviceability are just some of the things he makes a priority. That's why Forrest is excited about the new Bladeform 8100 Series blade server.

With support for multiple Quad-Core Intel® Xeon® Processors 5300 Series, the Bladeform 8100 Series takes advantage of Intel's proven reliability and energy efficiency. Forrest is very impressed that the Bladeform 8100 Series provides 90%+ high-efficiency redundant power supplies for operating cost reduction and earth-friendly computing. Up to ten blade servers, each with 2 removable hard drives, can be installed in the 7U chassis, allowing 60 total servers in a 42U rack. The Bladeform 8100 Series is a perfect choice for mission-critical enterprise applications, scale-out, and high performance computing environments.

When you partner with Silicon Mechanics, you get more than a powerful Intel-based Solution—you get an expert like Forrest.

intel Xeon® inside™
Quad-core. Unmatched.

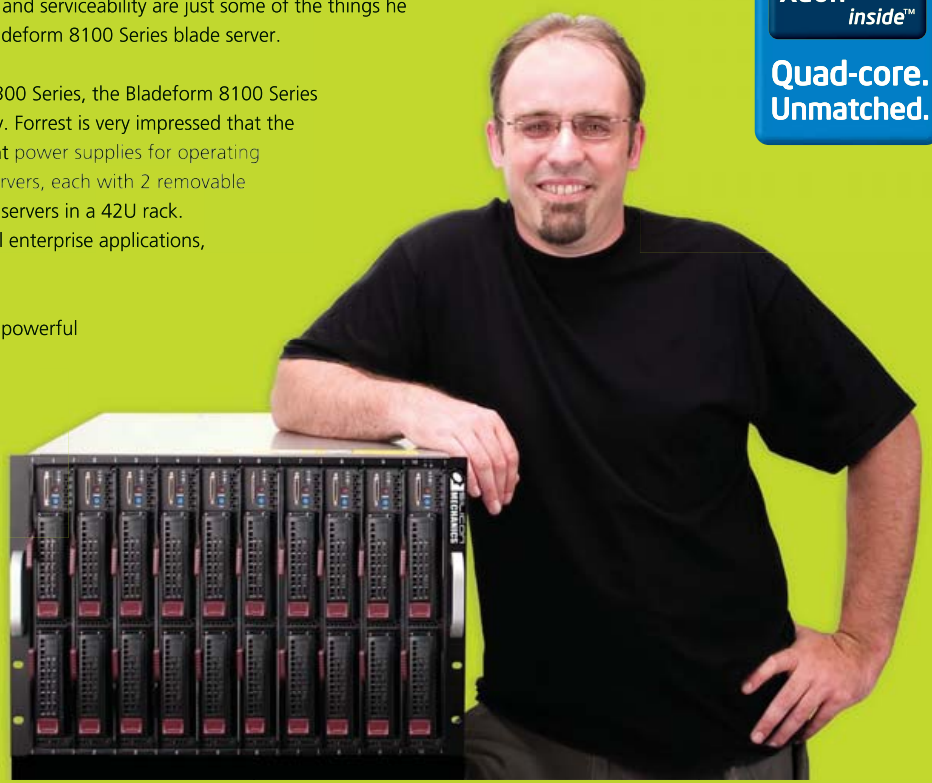

ATTENDING LINUX WORLD?
Visit Silicon Mechanics in Booth 414 for the chance to win one of three valuable raffle items.

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

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

of up to 550 miles per hour. Even when I zoomed out to approximate the view from my altitude, Navicore kept wanting to bring the view down to a few hundred feet over the Earth, giving me the sensation of NOE (nap-of-the-earth) flying in an F-16.

In case you're wondering if GPS use by passengers is okay, here is what it says in Continental's Flight Safety Regulations: "Devices such as electronic games, personal computers and entertainment players and recorders...must be used with headsets at all times. These devices, as well as noise-canceling headphones, calculators, shavers, cameras, GPS devices and aircraft power ports for laptops, may be used only at the gate when the main cabin door is open, or when announced by the flight attendants and the aircraft is above 10,000 feet in altitude. These devices must be turned off during taxi, takeoff and landing. Devices such as TVs, radio receivers and/or transmitters (including AM/FM/SW, CB and scanners), remote-control toys and personal air purifiers are not permitted for use at any time." My flight to Houston was on Continental. The flights from Santa Barbara to Los Angeles and Oakland were on United. I am a top-caste frequent flyer with United and have never been told not to use a GPS on a United flight at times when use of electronic devices are permitted.

One surprise for me was the performance of the little Nokia GPS unit. It was actually much faster at finding satellites than my Garmin. It worked at greater distances from the airplane window and saw more satellites as well. (On these flights, I put the receiver in the pocket closest to the window, and I had window seats on all the flights.)

The Navicore was a real winner in airplanes. Even though it didn't tell me my altitude, it was very handy for telling me where I was in relation to the ground below—a big plus for one of my favorite practices, which is shooting pictures out airplane windows. I have about 3,500 photos (**flickr.com/photos/docsearls/tags/aerial**) on Flickr tagged "aerial". With the excellent sensitivity of the little GPS receiver, the Navicore setup was actually better than my Garmin as an accessory to photography. The display of the N800 is also much larger than the Garmin's, which is a plus for my age-compromised eyesight.

Unfortunately, the Navicore kit was disappointing on the ground.

The main problem is speed. It's too slow at too many things. Finding locations and waiting for Navicore to create a route is often so time-consuming that one is inclined just to give up. Worse, once you find something, there is no obvious way to say "go there" and create a route. For that, you have to "Plan route" by connecting two waypoints.

For example, it took about a minute of poking and waiting to plan a route between the hotel in San Francisco and the Oakland International Airport. On a standard Garmin unit in a rental car, this takes only a few seconds. To be fair, the difference is partly due to the Garmin being a single-purpose device, while the N800 is an all-purpose computer. On the other hand, the Navicore system would benefit enormously from, say, predictive text. Most or all of the modern car navigation systems have this feature, and its absence with the Navicore is annoying. Also, switching between views with the Navicore can be torturously slow. When you tap on a screen icon for another view, 30 seconds or more might pass before anything happens. And, if you make the mistake of tapping on it more than once (because nothing happened the first time), the software might take what seems like forever before it finishes obeying your commands. Many times I have just given up and walked away until the system finally stopped thinking.

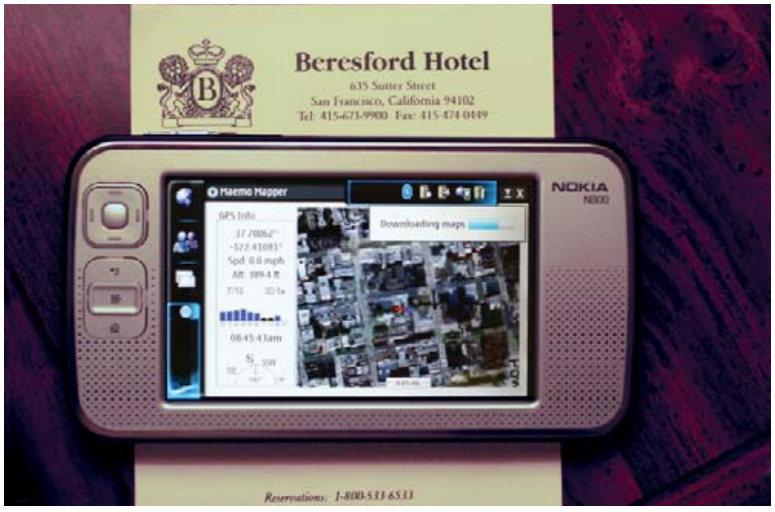

Maemo Mapper at the hotel in San Francisco. The GPS receiver was at the window on one side of the room. The N800 here was on a table across the room near the door. Over on another table was the Nokia E62 mobile phone, which the N800 used for downloading maps off the Web. All were connected by Bluetooth.

Here's what Maemo Mapper showed after the N800 and the GPS receiver were left on all night.

Three collaborating instruments provided by Nokia: the N800, an E62 phone and a GPS receiver. The phone used the Cingular/AT&T network.

The N800 could see several Wi-Fi access points in addition to the cellular connection, and showed whether the connection was locked.

Here's the GPS data page in Navicore. The other two tabs on the top are settings (the gear) and Bluetooth.

The Navicore's on-the-road navigation also isn't very smart. If you make a mistake on a Garmin (and I've had experience with a number of those), it says "recalculating", then changes the route you'll need to take. The Navicore system tells you to make your next available U-turn. I'm not sure if it wants you to backtrack to the departure point or to find an alternate route. Whether the logic is lazy, dumb or just wrong, the result on the road is annoying, useless or both.

Map orientation is another problem. If you want Navicore to be looking in the same direction that you're traveling (the forward view of the driver, in most cases), the system often switches for no apparent reason to "north up". Among lesser problems, the battery level for the GPS always showed 0%, even when fully charged, and clicking on the green satellite icon from the map view takes you to the control panel instead (where there's an icon for the satellite icon that does take you to the right place).

Maemo Mapper is a very different animal.

First, it's open-source freeware. Written in C by John Costigan and a crew of eight other listed developers, it is released under the GPL and is one of the most active projects at Maemo Garage (**garage.maemo.org**). I installed v1.4.6, which became available on June 9, 2007. After using the Navicore, it was gratifying to see the bug lists (15 open, 39 total), feature requests (33 open, 51 total), forums, project documentation, code repository and mailing list.

Second, Maemo Mapper is not meant strictly as a nav system, but rather as a multi-purpose GPS with a growing box of tools and uses. To explain the background and how it works, here's John Costigan, writing on his own Maemo Mapper site:

> The good thing about starting from scratch is that Maemo Mapper was built for the Maemo platform from the outset. No after-the-fact porting here—everything is 100% Hildon UI. Plus, I coded it with the form factor of the Nokia 770 in mind, so it is written in C with memory, CPU and screen-space optimizations in mind....
>
> Information about how to set up your maps directory structure is nice if you plan on lugging all of your maps around on an MMC card (as I do), but it's not for everyone, so there is an alternative: on-the-fly map downloading....
>
> Maemo Mapper was designed from scratch for the Hildon UI of the Nokia 770. The entire drawing area is devoted to the map, without cluttered GUI components. Most of the functionality (zooming, panning, opening route files) can be done using only the hardware buttons (i.e., without a stylus)....

All download operations are asynchronous and thus do not interrupt your user experience. Thus, you can use Maemo Mapper as you normally would while the maps are downloaded in the background, and the display is updated as the maps come in. This is especially noticeable when downloading maps automatically—the effect is similar to that of Google Maps or Google Earth....

Maemo Mapper supports opening and saving of standard GPS Exchange (GPX) files. You can open or save tracks (where you've been) or routes (where you would like to go). With this functionality, it is possible to download a route from the Internet (e.g., from the GPX Driving Directions Web service) and display the route on the map....

In addition to downloading third-party routes from the Internet, you can also generate driving directions from within Maemo Mapper directly. This functionality uses the aforementioned GPX Driving Directions Web service to generate directions to any location that would be recognized by Google Maps (e.g., "1250 Broadway St., SomeCity, PA" or "44012").

You can even specify that Maemo Mapper automatically update the route based on your current position, giving you the full power of a $2,000 in-car navigation system in the palm of your hand. Of course, an Internet connection is required during the trip....

Maemo Mapper will automatically notify you of upcoming waypoints. Waypoints are defined in Maemo Mapper as any track point in the GPX file that contains a description, and they are typically used in the context of Maemo Mapper to describe action that must be taken at the waypoint (e.g., "Turn left at Bourbon Street").

The notification takes the form of a Hildon "infoprint" notice (a pop-up in the upper-right corner of the display). Optionally, if you have Festival Lite (flite) installed, Maemo Mapper will use flite to synthesize a voice-based notification of the contents of the waypoint's description, thus providing you with an audio indicator of when and where to turn.

The amount of time/distance in advance that Maemo Mapper will announce the coming of a waypoint is configurable in the Settings dialog box. Also, the path of the flite binary to use is configurable in the Settings dialog box.

In the week since receiving the Navicore package (with the GPS receiver), I've installed and made use of everything I could that John talks about here. While I haven't been able to emulate a "$2,000 in-car-navigation system" yet, I have enjoyed using the N800 as a handheld GPS with the promise to rival many Garmin, Magellan or Tom-Tom units.

The Bluetooth connections between the N800, GPS receiver and cell phone (Nokia E62) are seamless and remarkably fast, flexible and error-free. I just took the N800 into another room while it downloaded maps and stayed in touch with the GPS receiver, even though both the phone and the receiver were up to 20 feet away (and apart from each other as well).

Maemo Mapper can use any number of maps. I have it set up to use these:

- Google Street.

PEG® the leading GUI for Embedded Systems

Prototype Today!!! FREE EVALUATION PACKAGE

- PEG+ - Full Featured Windowing in C++
- C/PEG - Smallest Footprint in ANSI C
- Royalty Free
- Fast execution speed
- Completely ROM-able
- Delivered with Full Source Code
- Development Tools including FontCapture, PEG WindowBuilder, and ImageConvert
- Complete set of screen drivers included
- Completely customizable
- Industry leading RTOS Support
- Supports all popular target processors, video controllers and I/O devices
- Multi-lingual support – 2-byte character sets & UNICODE string encoding
- Event-driven programming model
- Application Design Services
- Knowledgeable and timely support to users around the globe
- Now includes a fully licensed version of Paint Shop Pro 9

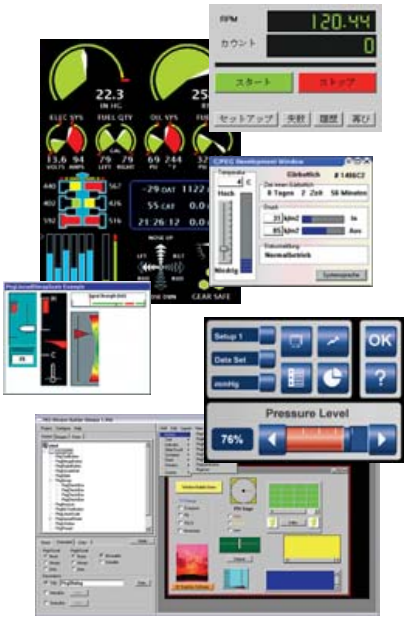

SWELL software

GRAPHICS SOFTWARE FOR EMBEDDED SYSTEMS

WWW.SWELLSOFTWARE.COM | 810-982-5955

- Google Satellite.
- VE (Virtual Earth) Street.
- VE Satellite.
- VE Hybrid.
- OpenStreet.

You can choose maps from the menu. Switching from one to another is quick and easy. You can zoom in and out with the N800's + and – buttons. These load very fast on a good Wi-Fi connection, and fast enough using the E62's GPRS connection through Cingular (now AT&T).

I'm looking forward to trying other mapping services as well, including TerraServer. Meanwhile, the current list gives me lots of angles to explore. Virtual Earth (from Microsoft) is especially nice, with a number of advantages in look and feel over the Google alternatives.

Navicore's map view zoomed out all the way, showing progress from Oakland to Houston.

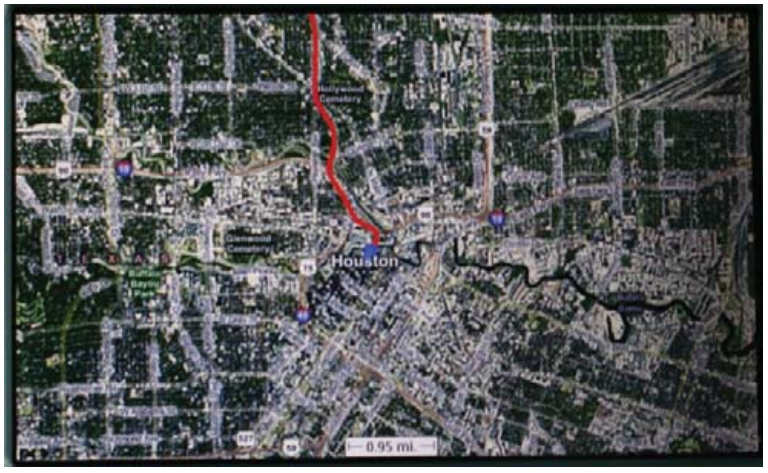

In Maemo Mapper, a red line with a blue pointer at the leading end draws the taxi's track from the airport toward our hotel in downtown Houston. Note how straight the red line is in this case.

Maemo Mapper provided a lot of fun while walking and taking taxis around San Francisco and Houston. Oddly, the red track left on maps in Maemo Mapper wanders off course at lower (walking and stuck-in-traffic) speeds. When idle, it might wander up to several blocks away from where it sits, leaving a bug-splat of red scribbles on the screen—not sure what's up with that. But, I am sure that, if it's a bug, the Maemo Mapper developers will fix it.

Flying was less fun with Maemo Mapper because it wanted to download map tiles, and there was no Net connection to allow that. (Though it was cool to see how many satellites I could get looking out one side of the plane and seeing how fast and high we were flying.) I'm not sure if there's a way yet to download and store maps for whole states or parts of continents, but again I have faith that these things will be worked out—if they aren't already—by the Maemo Mapper team.

In summary, it seems unfair to compare the two systems, because their main purposes are so different. But, the contrasts in speed and versatility between them are hard to ignore.

I wish I could recommend the Navicore system, because I really want the "N" platform to succeed—and I think it's great that Navicore is breaking the ice in what I hope will become a competitive market for software, hardware and services for handheld Linux devices. I also hope Navicore listens to our feedback (and other feedback as well) and continues to update and improve its offering.

Of course, I strongly recommend that Navicore embrace open-source development methods, in addition to embedded Linux hardware. Although there has been no official word from Garmin, I've heard from both Linux developers and Garmin staffers that future generations of Garmin gear will be built around Linux operating systems and sensibilities, and will no longer shackle users to Windows-only software for installing maps on attached GPS units. Rather, the units will simply be detachable storage devices, as OS-neutral as a USB thumbdrive. Navicore would be wise to take the same kind of approach. At the very least, opening up the development process would offer the best chance for continuing to work out the bugs that make Navicore's software so slow and frustrating. It would add value while not subtracting a penny from the price Navicore charges for the whole package.

As for Maemo Mapper, it's a great app that will only get better.■

Doc Searls is Senior Editor of *Linux Journal*. He is also a Visiting Scholar at the University of California at Santa Barbara and a Fellow with the Berkman Center for Internet and Society at Harvard University.

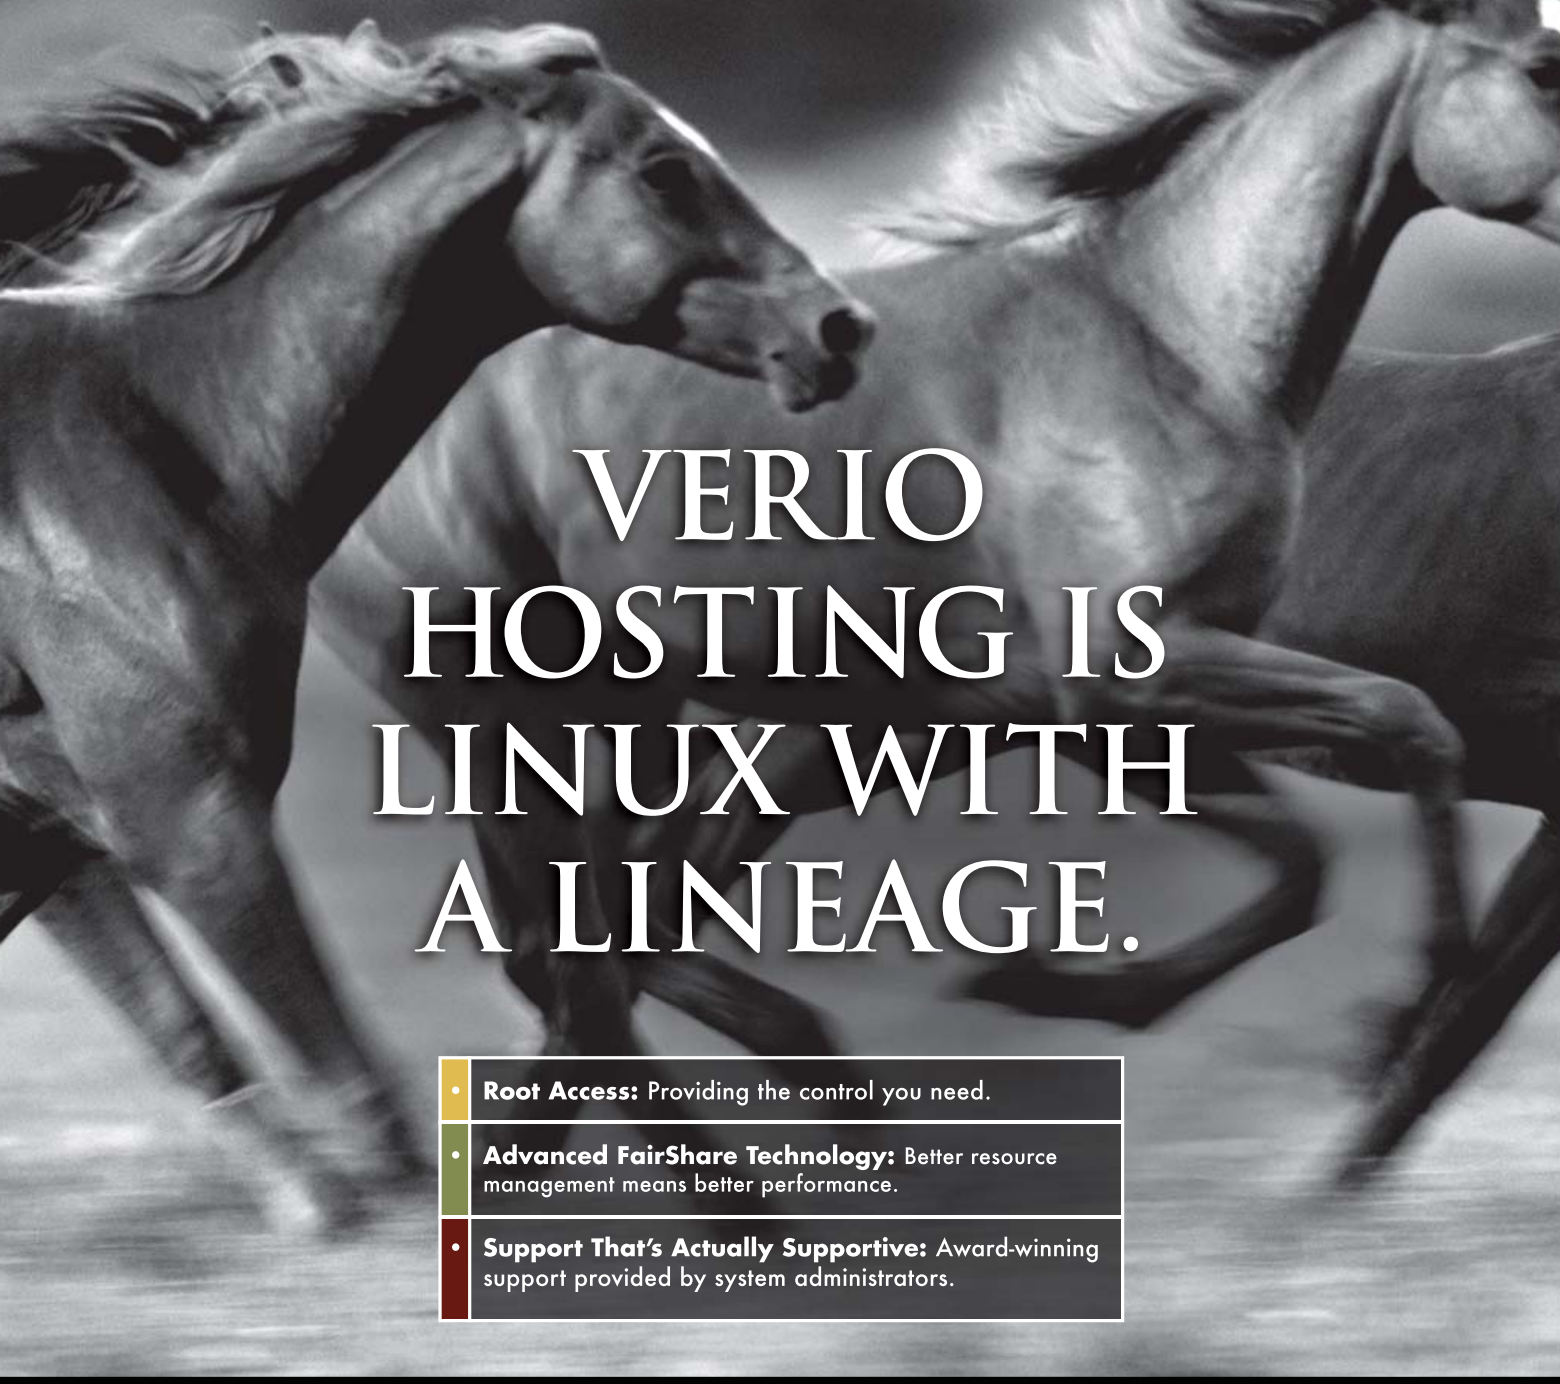

- **Root Access:** Providing the control you need.
- **Advanced FairShare Technology:** Better resource management means better performance.
- **Support That's Actually Supportive:** Award-winning support provided by system administrators.

Verio Linux® VPS and MPS: Best of Breed.

At Verio, we have a long-running commitment to open source, dating back to our early work with FreeBSD. Now, as the pioneer in virtual private server (VPS) technology and as a hosting provider backed by the financial resources of the world's largest telecommunications company, we bring something extra to Linux: reliability.

To learn more about Verio Linux VPS or Verio Linux MPS, or to register for a free test-drive, call 1-877-837-4654 or visit www.verio.com/linuxlineage.

Build on us.

VERIO

An NTT Communications Company

Verio and the Verio logo are trademarks and/or service marks of Verio Inc. in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. The mark FreeBSD is a registered trademark of The FreeBSD Foundation and is used by Verio Inc. with the permission of The FreeBSD Foundation. All other names are trademarks or registered marks of their respective owners. ©2007 Verio Inc. All rights reserved.

Beautiful Code: Leading Programmers Explain How They Think (O'Reilly Media)

Want to explore your inner Leonardo da Vinci? Then, pick up O'Reilly Media's new book *Beautiful Code: Leading Programmers Explain How They Think*, a compilation of thought-provoking essays that illuminate

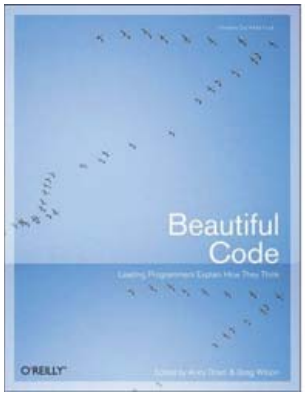

the artistry involved in crafting software. Edited by Andy Oram and Greg Wilson, *Beautiful Code* presents wide-ranging contributions from 38 pioneering software designers who aim to "rouse and inspire a new generation of coders" by sharing their most closely held secrets. For example, contributor Ronald Mak expounds on his elegant approach to NASA's Mars Rover, and Arun Mehta focuses on the importance of making technology useful to the disabled. Editors Oram and Wilson call the essays from such gifted inventors "inspiring and even uplifting". All royalties from *Beautiful Code* will be donated to Amnesty International.

www.oreilly.com

University of Maryland's Parallel Processing Chip

Researchers at University of Maryland's A. James Clark School of Engineering unveiled a new, yet-unnamed prototype computer that brings parallel processing to the desktop, increasing computing speeds up to 100-fold. The prototype, developed by Professor Uzi Vishkin and his team of Maryland Terrapins, utilizes a license-plate-sized circuit board mounted with 64 parallel processors. Further, Vishkin et al. developed the parallel-computer organization that allows the 64 processors to work together and make programming practical and intuitive for software developers. Although parallel processing has long been common in the supercomputing space, its application to desktop systems has been challenging due to programming complexities. However, the Vishkin team has made it possible via its novel single-chip parallel-processing technology. Future devices may include 1,000 processors on a single chip. Vishkin invites the public to name his technology, with a $500 cash prize going to the winner.

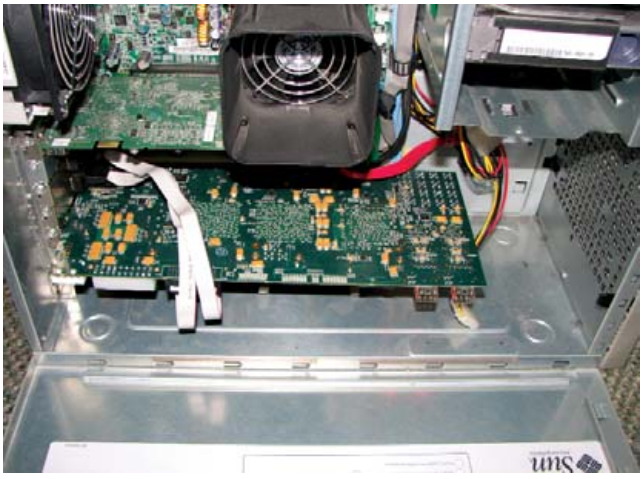

www.eng.umd.edu

Server Technology's Sentry Modular Cabinet Power Distribution Unit

The power requirements in today's datacenter have sparked innovations like Server Technology's new Sentry Modular Cabinet Power Distribution Unit (CDU), a family of modular solutions that enable simple, targeted distribution of 3-phase power. Sentry CDUs take a 3-phase feed-in and distribute single-phase power to any point in the cabinet. A single in-feed of 3-phase 208V 30/60A or 400V 32A can provide from 10.8 to 22.2kW of power. For greater densities or for redundancy, dual in-feeds can be provided with two modular units linked together under one IP address, providing anywhere from 21.6 to 44.4kW of power. IP access and security include Web interface, SSL, SSH, Telnet, SNMP, FTP, SNTP, Syslog, LDAP, LDAPS, TACACS+ and serial RS-232 access.

www.servertech.com

Paradigma Software's Valentina Community Database Server

We at *LJ* relish the torrent of ports to Linux, and we extend a warm welcome to a new arrival, Paradigma Software's Valentina Community Server 3.1 (VCS). VCS is a free and reportedly "very fast" Linux-based SQL database server whose kernel architecture is built for speed and can be optimized to allow "fewer and much smaller tables compared to other servers", says Paradigma. In addition to standard features, Valentina offers link refactoring commands, improved binary links (over ObjectPtrs) and object-relational design that "doesn't force one into a particular design pattern". Support for PHP 4.0 and higher is included, and Ruby support is on its way. Mac OS X and Windows versions also are available.

www.paradigmasoft.com

Cepoint's iS900 NAS Storage Time Server

The new kid on the time-server block is Cepoint Networks' iS900 iSCSI NAS storage/time server system with GPS or atomic time synchronization features. Expandable to 24TB, the iS900 is designed to enable synchronized, precise backup timing for organizations with geographically dispersed storage locations that must be backed up at a precise time. In addition to the time-synchronization functions, the iS900 supports RAID levels 0, 1, 10, 5 and 50 and various built-in encrypted remote data transfer features. Other features include continuous data availability, clustering, remote data replication, backup and bare-metal recovery, remote boot over an IP network and more.

www.cepoint.net

SourceKibitzer's Bio Service for Java Developers

SourceKibitzer, an established resource for Java developers working in open source, recently added a free Bio Service for enhancing professional opportunities. The Bio Service creates an on-line "resume" that includes 11 objective achievement metrics related to both the developer and software created. These metrics allow developers to market their contributions and real-world skills accurately to the Open Source community. The service also includes the ability to network within a larger community where members learn, interact, collaborate and compare their skills and experiences with one another.

www.sourcekibitzer.org

CentricCRM's Centric Team Elements

A preview edition (Version 0.9) of CentricCRM's forthcoming Centric Team Elements suite is now available for public consumption. Centric Team Elements is an open-source, Java-based Enterprise 2.0 application—that is, one that unifies Web 2.0 tools such as wikis, blogs and RSS feeds into one unified, relational database-driven application. For instance, a company department can set up wikis and blogs, track project plans and establish RSS feeds with departmental news—and, all information is indexed and searchable. Centric Team Elements also shares the same architecture as CentricCRM, an open-source CRM application, allowing seamless integration with a company's front-office operations. The 1.0 production release is slated for Q3 2007.

www.centriccrm.com

Please send information about releases of Linux-related products to James Gray at newproducts@linuxjournal.com or New Products c/o *Linux Journal*, 1752 NW Market Street, #200, Seattle, WA 98107. Submissions are edited for length and content.

ULTIMATE LINUX BOXES

Handheld p.51
Laptop p.54
Box p.58

The Ultimate Linux Handheld

The Nokia N800 starts off the N-series of Linux handhelds with an indisputable winner.

DOC SEARLS AND JIM THOMPSON

Last year's winner in this category, the Nokia 770, has a younger sibling, and, as oft happens, the kid takes the cake. Nokia's N800, the follow-up to the 770, is smaller, lighter, better-looking, faster and has a larger brain.

Although the N800 bears a high resemblance to its older sibling and retains the overall layout, excellent 802.11 and Bluetooth radios and razor-sharp 4.1", 800x480, 225 pixel/inch color touchscreen of the 770, many things have changed, including:

- Built-in stereo speakers, instead of a single speaker.
- Better microphone positioning.
- The addition of a 640x480 Webcam.
- The addition of an FM receiver.
- An upgrade to USB 2.0 connectivity.
- Two full-size SD card slots with a supported capacity of 2GB each, (in testing, a 4GB card worked).
- Instead of the 770's single RS-MMC slot, a faster CPU (TI OMAP 2420 at 330MHz vs. the 770's OMAP 1710 at 220MHz), yielding about twice the performance of the 770.
- Twice the RAM (128MB).
- Four times the internal Flash (256MB).

Overall, the appearance of the N800 is fantastic. It looks like a finished product, while the 770 had a "prototype" look. This attention to detail is obvious—from the newly designed aluminum front cover down to the built-in stand that locks at both 45 and 90 degrees. It even includes the little touches of chrome and lettering on the slide-out stylus.

Yet, even with all the additional functionality, the N800 is smaller (5.7" x 3.0" x 0.5") and lighter (7.3 ounces) than the 770 (5.5" x 3.1" x 0.7" and 8.1 ounces). The VGA resolution Webcam is accessed via a pop-out button on the left side of the screen.

One feature we miss is the reversible hard-shell cover, which served as a screen protector when closed. If Nokia (or someone else) offered one as an accessory, we would buy it. In addition, we found the smaller buttons on the top somewhat more difficult to manipulate.

The N880's software also is upgraded from the 770's. The company continues to ship a Debian-based distribution, with Hildon widgets, now updated as Internet Tablet Edition 2007. Improvements include:

- Updated Opera 8 browser with Flash 7 support.
- Jabber-based IM, including VoIP and video conferencing.
- Improvements to the connectivity manager so it remembers your connectivity preferences.
- The ability to use part of the internal SD card as swap, increasing virtual memory.

Nokia has announced (via the Maemo Project) a road map, including GTK 2.10, Samba, Bluetooth headset support, USB host support, Skype and improvements to the built-in e-mail client.

We continue to be impressed by the range of the N800's Wi-Fi chipset. It nearly always picks up APs (access points) that most notebook computers fail even to see. When a Wi-Fi connection is available, you can use the Bluetooth radio to gain access to the Internet via a cell phone. When we received the Nokia N800 to review, it included settings for connecting to Cingular. One of us (Doc) has a Nokia E62 mobile phone with a Cingular data plan, and his N800 review unit got on the Net with no trouble at all.

But, your mileage may vary. To perform the required setup manually, select Tools→Control Panel from the Application menu, navigate to Connectivity, double-tap it, edit the connection titled Cingular Internet, tap Next, and then enter the following connection information:

- Access Point Name: wap.cingular.
- Dial-up Number: *99***1# (this is unchanged).
- User Name: WAP.CINGULARGPRS.COM.
- Password: CINGULAR1 (also unchanged).
- Finally, tap Finished.

If you're a US T-Mobile customer, the access point name is one of these: internet.voicestream.com, internet2.voicestream.com or internet3.voicestream.com, depending on your plan and region. The user name and password fields are blank for T-Mobile.

After you get this set up, power-cycle your phone, and you should be connected. Some phones (such as the E62) want you to approve making the connection. After clicking yes on the phone, you're on the Net. Although connectivity over GPRS isn't as fast as Wi-Fi, it's still a lot of fun to be able to pop up your favorite Web site at the beach. (And, both of your authors live by beaches—Jim in Hawaii and Doc in Santa Barbara. Your recreational options may vary.)

Unannounced—but recently discovered by the Nokia N800 community—the device has a built-in FM stereo tuner. Enabling this is as simple as updating the apt-based application catalog and searching for the FM receiver tool. Download it, and you're set once you plug in the headphones for use as an antenna. (FM waves are much longer than cell-phone waves, so you need a conductor up to 30" long or so.) You can configure the FM widget to play through the speakers, though this setting is not retained in the version we tested. (It defaults to the earphones.) Reception is about the same as you'd expect from a Walkman-type radio. One obvious omission we would like to see corrected: the NXP (née, Philips) TEA5761 FM receiver could be upgraded to a TEA5764, which includes RDS (Radio Data Service). That's the data stream that runs in the background with many US FM stations as well as all stations in Europe and many others around the world. For the N800, it would allow the FM tuner to identify stations, programs and music information. (Right now, on the FM tuner widget, you have to enter call letters manually for all your presets.) More important, RDS support would allow applications to be built that make use of data on the RDS stream. These include song or program title (date and timestamped) and other identifying information. Through ProjectVRM (which Doc heads at Harvard's Berkman Center), discussions already have begun toward making the N800 an interactive FM/streaming/podcasting device—one that could include a "buy button" that allows listeners instantly to become contributors to (and members of) public radio stations playing programs that listeners like. This opt-in orientation toward paying for radio (and podcasting) has the potential to revolutionize the whole business.

We even could use opt-in payments to fix commercial music radio. The Copyright Royalty Board earlier this year came up with a ruling that said every station on the Web would have to pay a tiny per-person/per-song fee (which increases to $.0019 each by 2010). While the Royalty Board and its predecessors (going back to the DMCA in 1998) spoke about a "hypothetical" marketplace with a "willing buyer and willing seller", the N800 actually makes such a marketplace possible, but on a completely voluntary, natural basis. Listeners pay only for what they like, in any amount they like, whenever they like. They

can subscribe, escrow the data for use (including payment) at a later time or whatever. The key is providing a low-friction way of listening and paying without any coercion on the supply side—a way to blow up DRM by turning stations and podcasters into true intermediaries between performing artists and those who enjoy their work. And, wouldn't you rather pay directly for public radio than have to listen to pledge drives twice a year?

Even without that extra feature, the N800 is a big advance on radio-as-usual. For example, you can "time-shift" radio, like TiVo does for video.

The N800 also has both USB and line-in interfaces, allowing you to record (and perhaps even edit) as well as play back podcasts. Simply add software. That's a huge advantage of the N800's nature as a wide-open Linux device.

One more fun thing (among many) to do with your N800 is use it as a navigation device. Nokia sells a Navigation Kit that includes a Nokia LD-3W Bluetooth GPS, car charger, car mounting kit, 2GB memory card and Navicore software with maps. However, if you already have a Bluetooth GPS, you might investigate Mameo Mapper. After you get through the setup, you can end up with a set of maps that are based on Google Maps (street or satellite view) or any of several other on-line map sources, including Microsoft's excellent Virtual Earth (VE). If you enable auto-download, these can be downloaded as needed through your cellular-phone data connection. (For more on Navicore and Maemo Mapper, see Doc's Linux for Suits column in this issue.)

We found that we could get five hours of use on a single charge while doing typical Web browsing, over Wi-Fi or Bluetooth, and listening to the FM radio. This is impressive for a device with this form factor and a Wi-Fi radio.

It's still early, but it's clear that Nokia is committed to the N-series (as it's now called) as a platform, and so are both vendors (such as Navicore) and developers. Be sure to visit both **maemo.org** and **garage.maemo.org**, where you can see what's developed and what's developing—and start rolling your own cool stuff for this excellent little device.■

Doc Searls is Senior Editor of *Linux Journal*. He is also a Visiting Scholar at the University of California at Santa Barbara and a Fellow with the Berkman Center for Internet and Society at Harvard University.

Jim Thompson has been noodling about with UNIX and Linux for far too long. He knows he started with BSD Unix Release 4.1a on a Vax 11/780 in 1980, and still thinks `echo 'This is not a pipe.' | cat - > /dev/tty` **is funny.**

Your World Runs Faster With c-tree® Database Technology

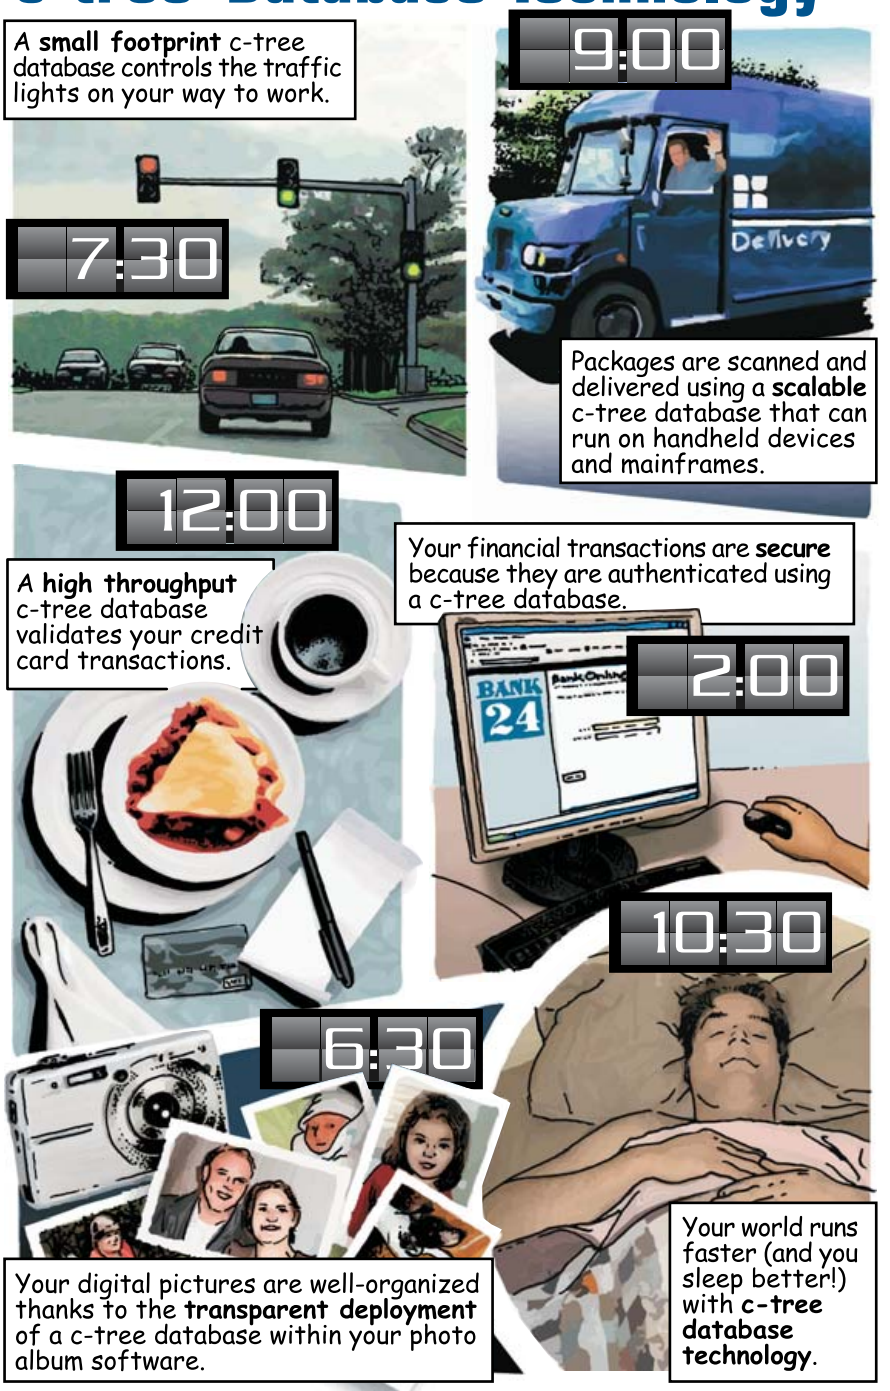

FairCom provides high performance, low maintenance data management technology. Our customers – ranging from small startups to multinational corporations – are able to overcome application-specific performance dilemmas because c-tree gives them precise control over their database operations. Super-charge your application and simplify your deployment! Download an evaluation edition of c-tree today.

FairCom®

Download FairCom's c-tree Plus Today!

www.faircom.com/go/?usesDLD

High Performance Database Technology • 800-234-8180

Other company and product names are registered trademarks or trademarks of their respective owners. © 2007 FairCom Corporation

The Ultimate Linux Laptop

EmperorLinux's Raven X60, take your victory lap! Smart innovations vault a compact yet powerful machine over the bar. **JAMES GRAY**

So, what does the word ultimate mean in the context of Linux-based laptop computers? Should it be the laptop to blow away all others? Or might the best model be more subtle in its "ultimateness"? Given that we use our laptops in such varied ways, the answer in this context is more elusive than on the desktop or in the server room.

Brainstorming for inspiration on the criteria for the Ultimate Laptop, and given that my laptop and I reside in Michigan, my mind wandered instinctively to cars. "Which vehicle would most people agree is ultimate?", I pondered. Based on what I see on local streets, I suspect that the Escalade, Corvette or XLR Roadster would win bragging rights. But, that's Michigan. An informal poll of my California friends gave me arguments for several different BMWs, Lexuses (Lexi?), Mini Coopers, Mustangs and more. What is a poor Products Editor to do?

Later, after reviewing the several entries for the Ultimate Linux Laptop competition, the ultimate automotive metaphor finally came to me. Although I probably will be run out of auto country for saying this, the laptop that was most like the Toyota Prius was the winner. In other words, the top machine was the one that would get *LJ*'s readers' heads nodding and saliva flowing with its novel and elegant—but not necessarily just brawny—Linux-based innovations.

Based on these criteria, the winner of the 2007 Ultimate Linux Laptop competition is the EmperorLinux Raven X60 Tablet. The Raven X60 is a portable, powerful, well-designed, Core 2 Duo machine with several unique innovations that Emperor has brought to the Linux space. Shortly, we'll delve into why we selected the Raven X60, but first, here's a little more about the competition itself.

About the Competition

For this first iteration of Ultimate Linux Laptop, we emphasized the innovation-focused "Prius effect" as mentioned above, considering the entire computing experience. We called on vendors to send us their best publicly available total package for around $3,000 US or less. Linux should be pre-installed, and all major features (such as Wi-Fi support, Bluetooth, function keys and so on) should work without major extra input from the user. Furthermore, after testing all the machines, we were more convinced that the complete package is much greater than the sum of its parts. Some machines looked wonderful on paper but simply failed to inspire. Factors such as ergonomic construction, service, documentation and attention to detail noticeably and considerably separated the best from the adequate.

We anticipate that some readers will prefer that the Ultimate Linux Laptop act more like a Porsche than a Prius and provide the most smoke-producing performance possible. Yes, we considered this, and in the future, we may bump up the price point to include exotic machines. This year, however, we felt that the current market for Linux laptops warranted a brains-over-brawn approach.

Luckily, many of the firms that ship pre-installed Linux machines took part in our competition. We cheered the news of Dell's shipping pre-installed Linux laptops and tried to get our hands on one for testing. Unfortunately, however, Dell expressed zero interest in participating, despite repeated calls around Round Rock. In contrast, I spoke with either the CEOs or Vice Presidents of all of the Linux specialist providers. Although many in our community make calls to support Dell, I say that actions speak louder than words. In my experience, our community-based providers generally do an excellent job of supporting what they sell.

Now, on to the Raven X60 Tablet laptop!

Figure 1. The Raven X60 has a 4,096dpi Wacom Tablet built in to the LCD display.

The X60 Tablet Stands Out

Here are the specs for the Raven X60 Tablet we tested, priced at $2,950 US:

- Dual-boot Ubuntu Feisty Fawn (kernel 2.6.21) and Fedora Core 6 (kernel 2.6.17), both 32-bit.
- Processor: Intel 1,667MHz Core 2 Duo L2400 with 2,048KB cache (667MHz FSB).
- 2,048MB of RAM.
- 100GB hard disk at 5,400RPM.
- DVD+/-RW (in UltraBase, included).
- 12.1" LCD display at 1400x1050.
- Networking: 10/100/1000Mbps Ethernet, 56Kbps modem, 802.11a/b/g (54Mbps) Wi-Fi, Bluetooth (2Mbps).
- Weight: 3.4 lbs. alone, 4.4 lbs. with battery, 6.4 lbs. with UltraBase.
- Battery: 8-cell, 4.5 amp hours.
- Ports: two USB, one CardBus, one RJ-45, one VGA, one secure digital.
- Bluetooth support.
- Wacom tablet integrated into LCD display.
- Biometric fingerprint scanner.
- UltraBase station includes DVD+/-RW, additional ports—parallel, serial, USB (four), PS/2—and pass-through of RJ-45, VGA.

Four features on the Raven X60 Tablet put it all alone into this year's winner's circle. First, the X60, which is simply a Linux-powered Lenovo ThinkPad X60 Tablet, is a well-built ergonomically efficient machine. Lenovo has continued IBM's strong tradition of quality, which is evident in the components used, the construction (such as the display swivel) and the comfortable computing experience (for example, the keyboard and stylus response).

The second stand-out feature is the 4,096dpi Wacom Tablet built in to the LCD display, which works like a charm. Emperor conveniently pre-installs the Jarnal handwriting recognition software with a desktop icon, so you can rotate and flip the display any way you wish, even flat on its back, and start writing. The Dynamic X rotation allows you to cycle between four screen orientations without restarting X. In addition, the included

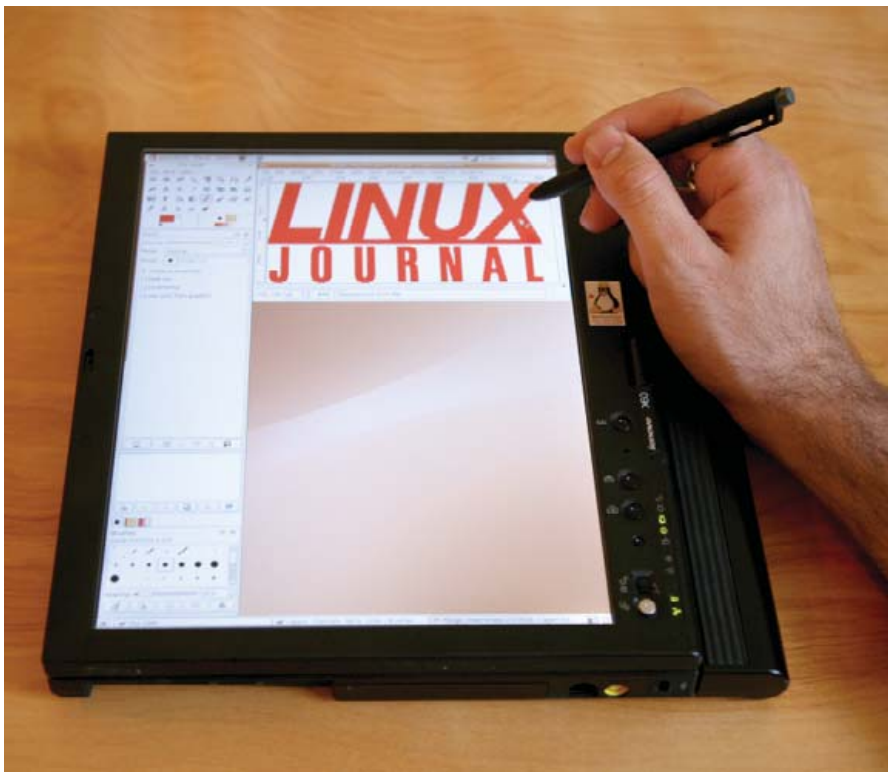

Figure 2. Use the tablet for your graphics work on The GIMP, which recognizes 256 levels of pressure. You also can cycle between four different screen orientations.

Linux Laptops
Starting at $799

Linux Desktops
Starting at $375

Linux Servers
Starting at $899

DON'T BE SQUARE!
GET CUBED!

R3 Technologies
309.34.CUBED
shoprcubed.com

Figure 3. The included UltraBase docking station contains the DVD+/-RW drive, as well as a multitude of additional ports.

stylus can act as a mouse or a stylus with a pressure-sensitive eraser. Besides tinkering with Jarnal, I had a blast playing with the tablet function on The GIMP, which made me feel like a budding artist. Emperor's president, Lincoln Durey, told me that The GIMP "takes advantage of the 256 levels of pressure sensitivity that the stylus provides".

Yet another slick feature is the Raven X60's biometric fingerprint scanner, which currently works on the GDM, login and sudo PAM services. Once you've trained your own fingerprints, you can use it to log in or simply revert to your password.

The fourth feature to distinguish this machine is its comprehensive, well-designed manual, which is specific to the Raven X60. It's obvious that Emperor considers good documentation a core part of the user's experience, as the manual nicely predicts what information the user will seek. Furthermore, each section clearly explains not only what features exist and how to use them, but also what features are not working yet. This service will save you hours by steering you away from dead ends.

The Frosting on the Cake

Besides the four aforementioned features, there is much else to like about the Raven X60. Regarding the things under Emperor's control, it provides a well-planned Linux experience. The Raven X60 came dual-boot with Ubuntu Feisty Fawn or Fedora Core 6 installed, both of which had many nice features. (However, Ubuntu was a little more polished than Fedora.) Standard features were all working out of the box (including wireless networking and Bluetooth), and with a few clicks on both distros, I was able to install the codecs that allowed me to play unencrypted DVDs, MP3s and Windows Media and Audio files.

The Raven X60 is a portable, powerful, well-designed, Core 2 Duo machine with several unique innovations that Emperor has brought to the Linux space.

Of course, one of the critical tests of a good pre-installed Linux laptop is how well the function keys work, and the Raven X60 performed well here. Despite sluggish brightness controls, functions such as volume control, external VGA and hibernation worked well. The media player controls worked, but only with Totem.

Emperor also ensured that the docking station, otherwise known as the UltraBase, works; however, it's not hot-plug yet due to it containing an IDE device, but Emperor says that this feature is on the way. The UltraBase, which ships standard with the Raven X60, holds the DVD+/-RW out of the box, but one also could swap in an extra hard drive if desired. It also has several additional ports (parallel, serial, four USB, PS/2) and pass-through for RJ-45 and VGA.

Other features where credit is due to Lenovo (and Intel) include the 1,667MHz Core 2 Duo processor, conveniently recessed ports and a non-wimpy display swivel that looks and feels like it will go the distance.

The only complaints to record are that the Raven X60 I reviewed came with 32-bit Linux rather than 64-bit, though Emperor indicated that, at the time of this writing, it is nearly done with tablet support for the latter. By the time you read this, you should be able to utilize Emperor's Web-based configurator to install 64-bit Ubuntu or Fedora for this machine. In addition, the hard drive I received runs at only 5,400RPM, and given that this was the Ultimate Linux Laptop competition, I would've expected to receive the faster model at 7,200RPM.

A Victory Lap for the Raven X60

Because Emperor started with a well-built Lenovo machine and then made sure that its bells and whistles all worked under Linux—most notably the tablet support, biometric fingerprint scanner and excellent documentation—it deserves to take a proud victory lap for its hard work on the Raven X60. These aspects show that the company clearly values attention to detail and wants to maximize your Linux-based laptop experience.

After the victory lap, however, we hope Emperor will continue to innovate and push forward the dynamic domain of Linux-based laptops. It will have to keep looking over its shoulder though, because the market for Linux laptops is maturing, the competition is fierce, and the market for exotics will get more interesting. See you next year!■

James Gray is *Linux Journal* Products Editor and a graduate student in environmental science and management at Michigan State University. A Linux enthusiast since Slack 1.0 in 1993, he currently lives in Lansing, Michigan, with his wife and kitty.

Polywell's Ultimate Linux Systems

More Choices, Excellent Support Service, Great Value!

1U Value Servers Starts at $399, Up to 4GB RAM

(custom config. available)

$399 1U-485Ax Sempron 3000+, 1GB DDR2, 80GB HD
(For Volume Purchase Only)

$499 1U-485Ax Athlon64 3500+, 2GB DDR2, 80GB HD
(For Volume Purchase Only)

$699 1U-690GA Athlon64 X2 Dual-Core 3800+, 4GB DDR2, 2x80GB HD (Dual LAN +$45)

$1299 1U-690GA Athlon64 X2 Dual-Core 4200+, 8GB DDR2, 2x80GB HD (Dual GigaLAN +$45)

Linux Appliance Starts at $299

Poly 485Ax Sempron 3000+, 512M DDR2, ATI X1100 Graphics, 100Mbit LAN, DVD, 160G HD $299 (For Volume Purchase Only)

Poly 690GA Athlon64 X2 Dual-Core 3800+ 1G DDR2 ATI X1250 DVI+VGA, GigaLAN, DVD-RW, 500GB HD $499

(OEM /ODM Service Available)

1U Advanced Servers, Up to 64GB RAM, 4TB HD

$1,999 1U-690S4 Athlon64 X2 Dual-Core 500+, 8GB DDR2 2TB 4x500GB HD, Dual Gigabit LAN

$2,399 1U-1000SL Opteron 1210 Dual-Core, 8GB DDR2 ECC 2TB 4x500GB HD, Dual Gigabit LAN

$2,999 1U-2500A16 2 x Opteron 2212 Dual-Core, 16GB ECC DDR2, 1TB 4x250GB HD, Dual Gigabit LAN

$5,999 1U-2500A16 2 x Opteron 2216 Dual-Core, 32GB ECC DDR2, 2TB 4x500GB HD, Dual Gigabit LAN (Option: 64GB+4TB HD)

Low-Cost NAS Storage 2.0TB Starts at $999

$999 Netdisk 4000 2.0TB 4x500G

$1,550 Netdisk 6000 3TB 6x500G (2U)

1U Twin Servers, 1U 8-Way Quad Opteron

$4,999 1U-Twin 2 x 2 Dual-Core Processors, 2 x 8GB ECC DDR2, 2 x Dual 250GB HD, 2 x Dual Gigabit LAN

$5,999 1U-8415A 4 x Opteron 8212 Dual-Core, 16GB ECC DDR2, 1.5TB 3x500GB HD, Dual Gigabit LAN

$7,999 1U-8415SS 4 x Opteron 8212 Dual-Core, 32GB ECC DDR2, 2x74GB 15K RPM SAS HD, 3x Gigabit LAN

High-Density Multi-Processor Servers

1U 8-Way, 5U 16-Way Servers, Up to 128G RAM

$12,999 1U-8450SS 4 x Opteron 8212 Dual-Core, 64GB ECC DDR2, 2x74GB 15K RPM SAS HD, 3x Gigabit LAN

$39,999 1U-8450SS 4 x Opteron 8214 Dual-Core, 128GB ECC DDR2, 2x74GB 15K RPM SAS HD, 3x Gigabit LAN

$18,500 5U-8850T5U 8 x Opteron 8212 Dual-Core, 64GB RAM, 2TB 4x500GB HD, 3 x Gigabit LAN

$46,999 5U-8850T5U 8 x Opteron 8214 Dual-Core, 128GB RAM, 4TB 8x500GB HD, 3 x Gigabit LAN

6TB 2U Storage Server 2012SC-2055A

4TB 8x500G, Opteron $2,999
6TB 12x500G, Opteron $3,999

Blade Servers - 10 Dual or Quad Processors Blades

$13,999 8U 10 x 2055A Blades
10 x (Dual Opteron 2210 Dual-Core, 4GB RAM, 80G HD)

$24,999 8U 10 x 2500M Blades
10 x (Dual Opteron 2212 Dual-Core, 16GB RAM, 80G HD)

$36,999 8U 10 x 8450A Blades
10 x (Quad Opteron 8212 Dual-Core, 16GB RAM, 80G HD)

$66,999 8U 10 x 8450A Blades
10 x (Quad Opteron 2214 Dual-Core, 32GB RAM, 80G HD)

18TB 4U Storage Server 4024AIS-2500A16

12TB 24x500G, Opteron $7,500
18TB 24x750G, Opteron $11,999

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

Polywell OEM Services, Your Virtual Manufacturer

Prototype Development with Linux/FreeBSD Support
Small Scale to Mass Production Manufacturing
Fulfillment, Shipping and RMA Repairs

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

888.765.9686

www.Polywell.com/us/LJ

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

Opteron, Sempron 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.

POLYWELL

The Ultimate Linux Box

We packed unbelievable power in a tank case and added all the trimmings for less than $4,000. **NICHOLAS PETRELEY**

What kind of computer would you build if you decided to spoil yourself, but keep your spending below the big price differential that occurs when you buy the next best thing? We put together a box based on that very principle, and it packs enough power to last years into a tank of a case. Then, we trimmed it with a beautiful 24" display, phat speakers, keyboard and mouse. The final product comes in at $3,887 US, with lots of wiggle room to save money on the pieces of the system you may consider more than you need.

Heck, it's almost all more than you need, but if we picked anything less, it wouldn't be the Ultimate Linux Box, would it? If your significant other questions your wisdom, you always can fall back on the following speaker analogy: why buy speakers that have a frequency response outside the range of human hearing? These speakers are more likely to perform well within the range of human hearing, right? And, that's why you indulge in overkill for the Ultimate Linux Box. Maybe you won't push it to its limits, but it will perform better within the limits of your work habits, right? Let me know if your significant other buys the argument. (Don't mention that whatever you buy will be obsolete in six months.)

Don't worry if you get "no" for an answer, or if you're strapped for cash. We didn't forget you. We tossed together a Penultimate Linux Box that totals less than $2,000 US (see The Penultimate Linux Box sidebar). It still packs an amazing amount of power. In fact, even the Penultimate Linux Box leaves wiggle room for saving money if you can get by with a slower CPU or display card. As you shave off options, you also can shave the price off the case and power supply.

But, that's not why most of you are reading this, I hope. You want to drool, and we chose some awesome hardware to get those juices flowing. It all revolves around a stunning ASUS motherboard with an Intel Core 2 Quad processor and 4GB of RAM, coupled with a 3ware RAID controller and RAID cage with four 320GB drives. Add one of the latest, greatest display cards, and you've eliminated performance bottlenecks at every turn.

Keep in mind that all prices are as of the time of writing of this article. Prices drop quickly, so you may be able to purchase better hardware for the same price or the same hardware for less cash. Consider also that products get discontinued (our first display card was discontinued a week after we tried it), and that vendors slipstream changes into hardware. Slipstreamed changes mean you may not get exactly what we tried even if you buy the same make and model of any given piece of the box.

Now that you've been forewarned, read on for the details about our 50-gallon drum of butt-kicking hardware.

Figure 1. Ultimate Linux Box

CPU

We chose as our ultimate CPU the Intel Core 2 Quad Q6600, two dual-core CPUs combined into a single part. Each core runs at 2.4GHz. It isn't a true quad-core CPU, so it doesn't scale like a true four-CPU system, but we quickly fell madly in love with it anyway.

You're probably already asking yourself if you really need four cores. Most likely you don't. But, you may find that you appreciate all of them, depending on the kind of work you do. Linux is noticeably more responsive with the Quad processor than with the dual-core AMD configurations we tried, and it's a thrill to watch things compile with the command `make -j 5`. It is common practice to have the make command spawn a number of processes equal to the number of CPUs plus one. The `-j 5` switch spawns five

The **Ultimate** Linux Box

MOTHERBOARD

- ASUS Striker Extreme LGA 775 NVIDIA nForce 680i SLI ATX
- 6 x 3.0Gb/s SATA (internal)
- 2 x 3.0Gb/s SATA (external)
- 1333/1066/800MHz front side bus
- DDR2 800
- Maximum 8GB memory, dual-channel
- Three PCI Express slots, two of them for SLI
- Dual 10/100/1000Mbps LAN
- SupremeFX audio
- Onboard switches (CMOS clear, Power, Reset)

Price: about $330

CPU

- Intel Core 2 Quad Q6600 Kentsfield 2.4GHz
- Model BX80562Q6600
- Socket: LGA 775
- FSB: 1,066MHz
- L1 Cache: 64KB+64KB
- L2 Cache: 2 x 4MB
- MMX, SSE, SSE2, SSE3, SSE4, EM64T

Price: about $515

MEMORY

- Patriot eXtreme Performance 4GB (2 x 2GB modules) 240-pin DDR2 SDRAM
- Timing 5-5-5-12

Price: about $250

RAID

- 3ware 9650SE-4LPML PCI Express
- 4 SATA II controller card
- RAID: 0, 1, 5, 10, Single, JBOD

Price: about $335

- 3ware RDC-SATA Internal RAID Drive Cage
- Three cooling fans
- Four drives
- Requires three 5.25" drive form-factor slots

Price: about $200

NOTE: combination prices available for RAID card + cage: about $500

DISK DRIVES

- Western Digital WD3200AAKS 320GB 3.0Gbps drives x 4

Price each: about $90; total price: about $360

- NEC Black 16X DVD+R Burner

Price: about $30

DISPLAY CARD

- PNY VCG8800UXPB GeForce 8800Ultra
- Memory: 768MB 384-bit GDDR3
- PCI Express x16

Price: about $670

POWER SUPPLY

- Thermaltake W0106RU 700-Watt power supply
- Modular cables
- SLI certified

Price: about $170

CASE

- Cooler Master Stacker 830 ATX Full Tower
- Nine 5.25" drive bays
- Four 3.5" drive bays
- Front ports: USB, audio, IEEE 1394
- Front 120mm fan
- Rear 120mm fan
- Up to four 120mm fans in side panel

Price: about $250

MONITOR

- Acer AL2416WBsd 24" 5ms DVI Widescreen LCD
- Brightness: 400 cd/m2
- Contrast: 1000:1
- Colors: 16.7 million
- Pitch: 0.27mm
- Recommend resolution: 1920x1200
- Viewing angle: 160 degrees horiz and vert

Price: about $550

KEYBOARD

- Microsoft Natural 4000 Black Wired Keyboard

Price: about $52

MOUSE

- Logitech G5 Laser Mouse
- Wired connection
- User-chosen weights
- On-mouse adjustable sensitivity

Price: about $45

SPEAKERS

- Creative I-TRIGUE L3800 48 Watts 2.1 Speakers
- Frequency response: 30Hz ~ 20kHz
- Signal noise ratio: 80dB

Price: about $130

processes, one for each core, plus one.

Granted, the Intel Core 2 Quad is pricier than AMD dual-core CPUs. But, the $515 price tag is very reasonable for the kind of performance you get, especially when you consider that better Intel CPUs sell for almost twice the amount.

Speaking of AMD, it says a lot that we ended up going with an Intel Core 2 Quad for our ultimate system. We're passionate fans of the AMD FX-74 (socket F, 1207 FX) and AMD FX-62 (socket AM2) dual-core processors. These processors cook, and even the 3.0GHz per core FX-74 chip sells for $100 less than the Intel Core 2 Quad. We like AMD processors so much, we almost chose as our favorite the combination of an ASUS M2N32-WS Pro with an FX-62. (The M2N32-SLI, the board we chose for the Penultimate Linux Box, is almost the same board, but it lacks a PCI-X slot.) This combination will be more than enough power for most people, and it's hard to beat the price of the FX-62 at the time of this writing ($280). If you choose this combination and want to

The **Penultimate** Linux Box

MOTHERBOARD
ASUS M2N32-SLI (no PCI-X slot): $170

CPU
AMD FX-62: $280

DISPLAY
EVGA GeForce 7950GT: $200

MEMORY
CORSAIR XMS2 2GB: $160

DRIVE (320GB)
Western Digital WD3200AAKS: $90

DVD BURNER
NEC Black 16X DVD+R: $30

POWER SUPPLY
Thermaltake W0106RU Power: $170

CASE
Antec Nine Hundred case: $140

MONITOR
Acer AL2416: $550

KEYBOARD AND MOUSE
Microsoft Natural: $52
Logitech G5: $45

SPEAKERS
Creative SBS Vivid 80: $34
Total: $1,921

add a RAID card though, make sure you get one that works with PCI-X instead of PCI Express. The 3ware 9550SX-4LP works well with this motherboard.

Of course, AMD processors cook in terms of heat too. We tried a number of fancy third-party fans, but none of them cooled the FX-62 processor below 95° Fahrenheit at idle. Worse, when we combined the FX-62 with our original water-cooled video card, the water cooler fan had to do double duty as a CPU case fan. We had to replace the water cooler fan with one that's more powerful just to keep the CPU at a normal running temperature. In the end, everything ran fine, but the system was much more noisy than it had to be.

In contrast, the Intel Core 2 Quad runs very cool and quiet with the stock fan. The cores run at between 72° and 88° Fahrenheit, which is quite a bit cooler than the AMD FX-62, even with a great third-party fan.

If you really want to go over the top on performance, you can invest in an ASUS L1N64-SLI WS Dual Socket L motherboard and plug in two AMD FX-74 CPUs for four processors. We didn't try that combination, but it stands to reason that it should blow away the Intel Core 2 Quad. The processors themselves are faster, and having two sockets theoretically brings you closer to what you expect from four CPUs in scalability. However, this configuration represents a big jump in total price. The two processors are more expensive than a single Core 2 Quad, and you'll need an extremely hefty power supply (more than 1,000 Watts) to power the processors and video card. That's a lot of power, which means you'll also be generating a lot of heat. Think of it as a trade-off. You may pay through the nose for an ultimate system, but you can lower the house thermostat and throw away your stove.

Motherboard

We went with the ASUS Striker Extreme LGA 775 NVIDIA nForce 680i SLI ATX motherboard. With one possible exception common to most motherboards, the layout is excellent. It includes onboard power, reset and CMOS clear buttons. (In order to prevent you from accidentally clearing the CMOS, you have to change a jumper for the CMOS clear button to work.) There is a rear LCD panel with backlight, but we didn't find a need to use it. The motherboard includes excellent built-in HD audio via an add-in card located above the display card slot. It can run up to 1,333MHz for the front-side bus. The Core 2 Quad processor we chose needs only a 1,066MHz front-side bus. There are six internal SATA ports and two more external ports, none of which we actually used, as we went with the 3ware RAID card. It supports up to ten USB connections, four of which are connected to the rear connectors on the motherboard.

The one exception to the otherwise excellent layout is that, like many motherboards, the one place to insert a RAID card happens to be right next to the fan of a two-slot display card. This means the RAID card may block the airflow of the fan for the display card. We provide more details about this throughout the remainder of the article in the appropriate categories.

Figure 2. The power and reset switches along with USB ports and drive activity light are recessed into the top of the case.

Figure 3. Under the lip below the power switch are duplicate USB ports, an IEEE 1394 port and audio ports.

ASUS provides a handful of nice extras with the motherboard. It includes extra sensors that you can arrange to monitor the temperature of just about anything you like. It also includes a directional microphone you can place on top of the monitor.

ASUS includes a few pin adapters to make it easier to wire things to the motherboard. You attach things like the power switch, reset switch, power LED and hard drive LED to one of these adapters, and then plug the adapter onto the motherboard. This way, you can detach and re-attach these wires all at once, even while the motherboard is mounted in the case without having to use a flashlight and needle-nose pliers. ASUS also includes USB and IEEE 1394 pin adapters for those folks who buy a case or other attachment that does something silly like include individual wires for the USB or IEEE 1384 connection.

We don't know why you'd want to overclock a system like this, but ASUS makes it easy to do from the BIOS setup screens. If you push the system too far and it fails to reboot, press the reset or power switch again and it should recognize that there was a problem. It will reset the BIOS and allow you to try again.

Memory

We chose Patriot eXtreme Performance 4GB (2 x 2GB modules) 240-pin DDR2 SDRAM, which sells for about $250. Linux can use all 4GB of RAM whether you run a 32-bit x86 or 64-bit x86_64 kernel, but most people won't really need 4GB of RAM even with four CPU cores. If you want to cut back on the total price, you can try the CORSAIR XMS2 2GB (2 x 1GB modules) 240-pin DDR2 SDRAM, which runs about $160. The latency on the CORSAIR modules is actually better than the Patriot eXtreme. The CORSAIR timing is: 4-4-4-12, and the Patriot is 5-5-5-12. We've run benchmarks that show that lower latency helps performance, especially with AMD processors, but if you can tell the difference in actual everyday use, let us know. We can't.

Regardless of the memory modules you choose, make sure you insert them into the correct slots for dual-channel operation. In the case of the ASUS Striker Extreme, it is the first and third slot (every other memory slot, with matching colors).

RAID

The ASUS motherboard has what it calls integrated RAID onboard, as is true of most motherboards available today. This is really a misnomer. There is no hardware RAID controller on this or most other motherboards. The onboard RAID is really just a multichannel SATA controller. We configured Kubuntu 7.04 to run RAID 0+1 (also known as RAID 10) using four drives attached to the onboard SATA. It worked fine, but it was much more trouble than it was worth, so we do not recommend that approach. It is an especially bad idea if you intend to run more than one version of Linux on the same machine. It was hard enough going through the procedure once. We wouldn't want to repeat the process for every distribution we tried.

If you really want to use RAID, take our advice and buy a real RAID controller card. We chose the 3ware 9650SE-4LPML PCI Express controller. We connected four drives and configured them in RAID 10, which provides the best performance and safety at the cost of disk space. It stripes two sets of two drives, mirrored. The striping gives you the performance. The mirroring gives you safety, because you can replace a failed drive without losing any data. However, because two drives are redundant, you get half the disk space of your four drives. Our four 320GB drives gave us about 640GB of disk space.

If you really want more storage space than we created, you can buy larger drives or sacrifice some performance and configure

Never tried SlickEdit?
What's your excuse?

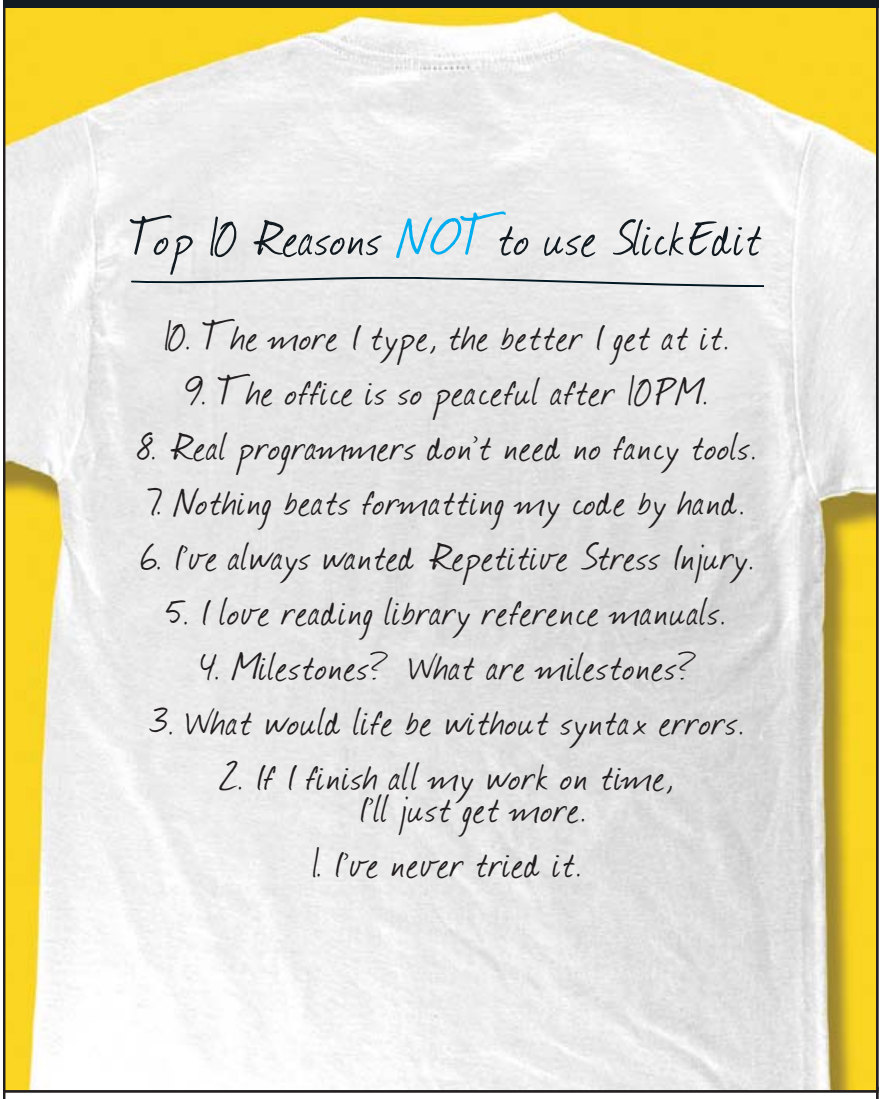

There's really only one good reason not to use SlickEdit — *you've never tried it.* How can we be so sure? User feedback. For 19 years we've listened to power programmers and continued to try to exceed their expectations. Our multi-language development tools and advanced code editors are fast, powerful, and flexible — whether you're developing on Windows, Linux, UNIX, or Mac OS. Whatever your environment, SlickEdit has a tool for you. Our product suite includes SlickEdit®, SlickEdit® Tools for Microsoft® Visual Studio® 2005, and the SlickEdit® Plug-in for Eclipse™.

slickedit

www.slickedit.com

Download a FREE trial at www.slickedit.com/trial or call 1-800-934-3348.

your array as RAID 5. RAID 5 trades write performance for more storage space.

The 3ware controller is superb. It delivers excellent performance and it is very easy to set up. You press Alt-3 to activate the setup screen at boot time. This utility allows you to create storage specifically for booting operating systems, but you don't need to use this feature. It may be necessary for other operating systems, but you can boot fine from the RAID array with Linux by using normal RAID partitions.

You can add a battery backup unit to the RAID controller so that you are less likely to lose data if you experience a power outage. We didn't include the battery backup as part of our ultimate box though.

You shouldn't need to add drives to your system via the onboard NVIDIA SATA controller if you use this RAID card. If you do add drives to the onboard SATA, however, be warned that some Linux distributions may get confused about the order of drives in your system. We tried adding a drive and did not experience this problem with this particular combination of components, but this problem has reared its ugly head with other similar configurations, so we assume it's still possible.

You may see an on-screen message at boot time that says the controller is not compatible with your BIOS. It goes by so quickly that you may miss it. If that concerns you, there are a variety of other RAID cards from which to choose. However, despite this warning, our 3ware card has performed without a hitch for weeks, and we love the fact that the Linux kernel has great support for the card by default.

The card is PCI Express x 4, which works fine in the middle PCI Express slot of the motherboard. If you go with our recommendation, make sure you plug the card in to the center PCI Express slot, not the second PCI Express slot for video SLI. You will experience lockups and problems if you choose the latter slot.

Figure 4. The 3ware RAID Cage with the Door Closed

Figure 5. Pop open the RAID cage door, unlock the drives and pull them out for easy hot-swapping.

RAID Cage

Some might consider a RAID cage to be frivolous. But this is the Ultimate Box after all, so we're including the 3ware RAID cage that lets you hot swap four drives in the space of three 5.25" drive bays. Aside from easy drive replacement, the RAID cage has two advantages over mounting the drives normally. First, the 3.5" hard drive cage in the case fits only two drives if you want good air circulation. You would have to mount at least two of the drives in the 5.25" area, and then add drive fans if you want to keep them cool. The RAID cage lets you pack all the drives in one small space, and it comes with its own set of fans to keep them cool. The cage also requires only two power sources, instead of one for each drive.

Display Card

We chose the PNY Technologies VCG8800UXPB GeForce 8800Ultra for our Ultimate Linux Box. This video card is one of the latest and greatest, which carries with it both advantages and disadvantages. The main advantage is clear: performance out the wazoo. Most people will never push the card to its limits. Despite the hefty price, you'd actually pay almost as much to plug in two cheaper cards in SLI mode, and you won't get nearly the same amount of performance.

On the other hand, the card is burdened with copy protection features meant for Microsoft Vista that you'll neither want nor use. It is built to support DirectX 10 as well. We may see DirectX 10 come to Linux sometime, but we're not there yet. Finally, if you pick any of the NVIDIA series 8 cards, such as this one, most if not all of the current Linux distributions will fail to detect it properly or set it up properly. You can configure the card to use the Xorg nv driver, or download and install the latest NVIDIA drivers yourself (see the How to Install NVIDIA Drivers on Ubuntu/Kubuntu sidebar for instructions). If you use the nv driver, there isn't much reason to go with a powerful card, because that driver doesn't make use of most of the power.

Our first pick was the NVIDIA-based MSI NX8800GTX, which sports its own water-cooling system. We chose an onboard water-cooled system because of the way many motherboards situate the display card next to a slot where you'd place a RAID card. The high-powered display cards take up two slots. The RAID controller card can block some of the airflow into the display card's onboard fan. You can put the display card in the second PCI-Express slot, but that usually interferes with PCI slot on the motherboard. Our configuration does not include a PCI card, so that may be a good option to keep the display card cool.

You won't have to deal with heating problems caused by an adjacent RAID card if you can situate a two-slot display card in the second PCI-Express slot, or if you buy our recommended video card or opt for one of the less-expensive one-slot GeForce 7 series cards. We tested a second sound card in our machine when we started. The sound card took up the PCI slot, which made it impossible to move the two-slot display card to the second PCI-Express slot. The onboard audio is great, though, so you won't have any problem using the second PCI-Express slot as long as you don't need some other PCI card.

If for any reason you do need to place a two-slot display card in the first PCI-Express slot, consider that RAID controllers tend to run a bit hot too, so this just adds to the problem. A water-cooled card like the MSI moves the fan off-board, which solves the problem. The solution worked beautifully with our Cooler Master case. You can remove the 120mm CPU case fan and mount the water cooler and fan in its place. This means the display card fan doubles as a CPU case fan. We had to replace the display card fan with a more powerful fan when we tried it with an AMD FX-62-based motherboard, because the AMD FX-62 runs so hot. In the end, either fan would work well with the Intel Core 2 Quad, which doesn't need as much cooling.

We would have kept the MSI card as our recommended display card except that it is no longer available. Perhaps that should tell us it has problems we haven't yet discovered. As it turns out, the PNY GeForce 8800Ultra is faster anyway. The position of the fan on the PNY card is such that the RAID controller card does not interfere with the airflow enough to cause any heating problems.

At $670, it's a very pricey card. We're

leaving it in as the default choice for the Ultimate Linux Box because it is pretty ultimate. We play games in our copious spare time (cough), so we like the way it handles 3-D graphics. Most games run—thanks to TransGaming's Cedega (although there are also native Linux 3-D games). Honestly, we're more likely to play around with 3-D rendering for amateur cartoons, so the rendering speed does come in handy.

Unless you do the same sort of thing, you won't need this much power. You can get a single NVIDIA 7950GT card instead, for example. The best of these cards generally runs at less than half the price of the 8800Ultra. Better yet, if you use an NVIDIA 7 series card like the 7950GT, you won't have to compile your own NVIDIA driver in most cases. It takes up only one slot, so you don't run into heating problems due to the proximity of the RAID card and the display card fan. You also can install two of these cards in SLI mode, which provides better performance without creating heat problems. However, two of these cards can cost almost as much as the 8800Ultra, and you won't get nearly as much performance for that money.

If you really don't need the best of the best in graphics, you still can get a screaming video card, such as the EVGA 256-P2-N636-AR GeForce 7950GT with 256MB of GDDR3 RAM for about $200, and there are plenty of decent lesser performers for less than $100.

Case

Our selection of cases boiled down to a contest between our favorite case from last year's Do-It-Yourself Ultimate Linux Box (the Silverstone TJO7-S full tower) and a new Cooler Master full tower case. We chose the Cooler Master case for both the right and wrong reasons.

Here are the right reasons. The case is almost tool-free. Almost everything snaps in and out without screws. We could jiggle in the RAID cage and flip a few plastic widgets to hold it in place, thus mounting the cage in seconds. The motherboard tray slides out, so you can mount the CPU, memory, video and RAID cards on the motherboard while the tray is outside the case. Slip the tray into the case, and it snaps into its proper place. The exceptions to the rule, where you'll need a screwdriver, are the power supply and possibly the hard drives. You'll need to use screws for hard drives only if you mount them in the 3.5" cage that comes with the case instead of inserting a RAID cage in the 5.25" bays.

The case is not without its drawbacks though. You should not have to do this, but we had to cut some plastic off the side panel in order to mount the water cooling fan for our first shot at a video card (Figure 6).

We had to make this modification only because the first video card we chose used a built-in water cooler with an external fan. The only reasonable place to mount the cooler was in place of the 120mm CPU case fan.

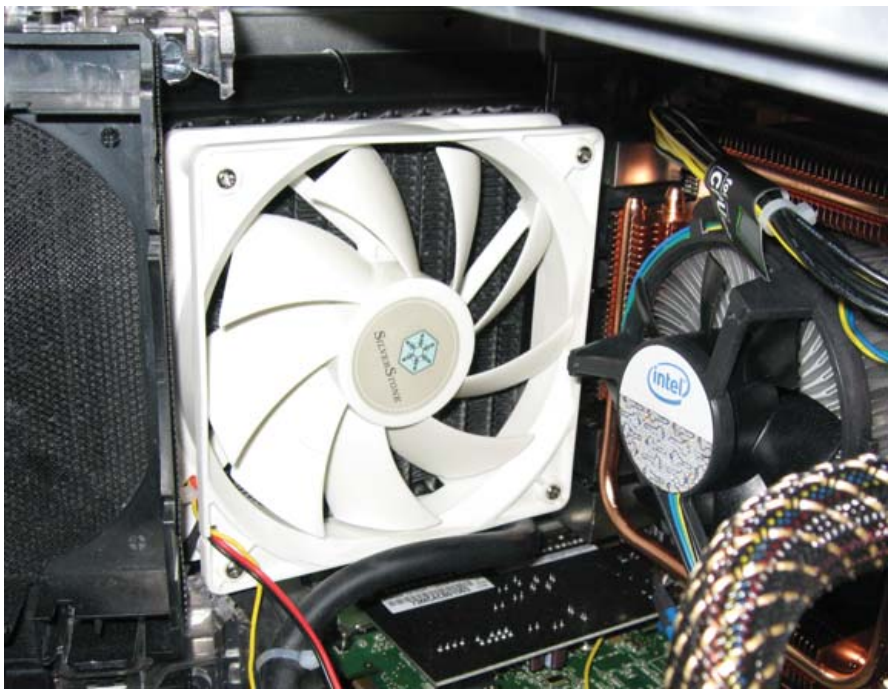

Figure 6. The original video card forced us to cut some plastic for the side panel door to close.

TS-7800 High-End Performance with Embedded Ruggedness

$269 qty 1 $229 qty 100

500 MHz ARM9

- New unbrickable design- 3x faster
- Backward compatible w/ TS-72xx
- Low power - 4W at 5V
- 128MB DDR RAM
- 512MB high-speed onboard Flash
- 12K LUT user-programmable FPGA
- Internal PCI Bus, PC/104 connector
- 2 USB 2.0 480 Mbps
- Gigabit ethernet
- 2 SD sockets
- 10 serial ports
- 110 GPIO
- 5 10-bit ADC
- 2 SATA ports
- Sleep mode uses 200 microamps
- Boots Linux in < 2 seconds
- Linux 2.6 and Debian by default

Design your solution with one of our engineers

- Over 20 years in business
- Never discontinued a product
- Engineers on Tech Support
- Open Source Vision
- Custom configurations and designs w/ excellent pricing and turn-around time
- Most products stocked and available for next day shipping

See our website for options, peripherals and x86 SBCs

Technologic SYSTEMS

We use our stuff.

visit our TS-7200 powered website at

www.embeddedARM.com

(480) 837-5200

How to Install NVIDIA Drivers on Ubuntu/Kubuntu

According to a recent survey, most *Linux Journal* readers use Ubuntu/Kubuntu, so here is one way to install the latest NVIDIA drivers on Ubuntu 7.04 and its spin-off distributions of the same version. (See our Tech Tips section for a more automated method.) Install the Linux source code, compiler and other build utilities first. You can use sudo for each command, but it's easier to get to a root prompt and work from there:

```
$ sudo -s -H
# (you should see this root prompt after you type the password)

# apt-get install build-essential linux-source
# cd /usr/src
# tar jxvf linux-source-2.6.20.tar.bz2
# cd linux-source-2.6.20
# make oldconfig
# make prepare
# make scripts
# cd /usr/src
```

Now, download the latest NVIDIA drivers, which we put in the directory /usr/src. *Important*: you must stop any graphical desktop you may have running in order to execute the NVIDIA installer. For example, if you are using KDM for graphical login:

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

Now, run the NVIDIA installer you downloaded. For example (assuming a 64-bit Linux installation):

```
# sh NVIDIA-Linux-x86_64-100.14.09-pkg2.run
```

Follow the installer prompts, and allow the installer to modify your /etc/X11/xorg.conf file.

You may have to make some changes to /etc/X11/xorg.conf in order to make X11 work with your card and monitor. The most important changes follow. Comment out the monitor's horizontal and vertical frequency range (the driver will discover your monitor's capability):

```
#    HorizSync       28.0 - 51.0
#    VertRefresh     43.0 - 60.0
```

Make sure the Screen section includes the maximum resolution you want to use. In the case of our choice of monitor, that will be 1920x1200 at 24-bit color. You can add other resolutions if you want to switch at runtime, but here's the bare minimum of what you want your Screen section to look like (various settings such as "nVidia Corporation" probably will be different on your system):

```
Section "Screen"
    Identifier     "Default Screen"
    Device         "nVidia Corporation"
    Monitor        "Generic Monitor"
    DefaultDepth    24
    SubSection     "Display"
        Depth       24
        Modes      "1920x1200"
    EndSubSection
EndSection
```

We even replaced the video card's fan with a more powerful fan because our first choice of motherboard used a hot AMD CPU. You could mount the cooler on the door, but the water tubes would prevent you from opening the door completely.

Regardless, even if the door interferes with anything, you don't have to cut any plastic off. As long as you don't install something that creates a lot of heat, there isn't much need to mount fans in the side door. You can remove the door entirely if it gets in the way of anything (it snaps out easily). The case has a side panel that normally fits over the door, so you won't even notice the door is gone once you have the case assembled.

It's a good thing that the Cooler Master case has casters to allow you to roll it around. This case is a tank. It is huge. It's even bigger than the Silverstone case, and we thought that was the biggest case we'd ever seen. Oddly enough, it doesn't feel like it's as big as the Silverstone case when you work inside it with the motherboard tray inserted. The difference is negligible, but we managed to scrape a knuckle now and then while arranging cables.

If we had this to do all over again, we'd try out the Antec Nine Hundred black steel ATX Mid Tower, which sells for only $140. We can't guarantee it would be a better case, but it's smaller than our choice, and the specifications still allow for plenty of room for the RAID cage and more. It even has the 120mm fan in the right place if we chose to keep our discontinued video card. The bottom line is that you won't really know whether you've got the best case unless you try them all, and that's clearly impossible. For what it's worth, we're very happy with the case we chose. But, we could have saved ourselves the aggravation of modifying the case if we'd held out a few more weeks before picking among video cards.

Figure 7. Add as many 4.5 gram weights to the puck as you like, and insert it into the mouse when you're done. This mouse is fully weighted.

Power Supply

The Thermaltake W0106RU 700-Watt power supply may sound like it is more than we needed, but it has a minimum output of 600 Watts, and a fully configured ASUS motherboard requires 550 Watts. We've had enough bad experiences with marginal performance power supplies that we're eager to err on the side of caution. Take heed that a single 8800Ultra display card requires two 12-volt connections. We can't imagine why you'd want to do this, but if you go crazy and install two of these display cards in SLI configuration, you'll need a better power supply than this one.

Monitor

We can't gush enough about the Acer AL2416 LCD monitor. Couple this puppy with a good display card, and you've got 24" of widescreen glory for only $550. You may be able to find a better display, but we couldn't find one that competes on all three categories of size, price and performance. The 1920x1200 display is crisp, bright, has sharp contrast, and it is fast enough even for gaming.

Speakers

The Creative I-TRIGUE L3800 2.1 speakers sound terrific, whether you're playing music or composing music with an attached synthesizer plugged in to the aux port. There's not much more to say other than pick whatever speakers suit your wants and needs. This is a simple 2.1 setup (stereo with subwoofer). If you have room and play games, go for the surround-sound speakers.

Keyboard

Microsoft should chuck its software business and sell keyboards. You can get much better ergonomic keyboards if you're willing to fork over lots of cash, but if you want to pay an average price for a keyboard, Microsoft ergonomic keyboards are among the best. We like ergonomic keyboards, so we went with the Microsoft Natural 4000 Black Wired Keyboard. Taste in keyboards is highly subjective, however, so just replace this choice with whatever you like best. It isn't likely to make much of a difference in the overall price.

Mouse

Taste in mice is almost as subjective as with keyboards, but you should give the Logitech G15 mouse a try. It has two killer features. First, it comes with a little puck in which you insert weights, after which you snap the puck into the mouse. This lets you make the weight of the mouse fit your personal preference. You might not think that makes much difference until you try it.

Second, you can press the minus button on the mouse to shift into multiple lower resolutions, and jump back to higher resolutions by pressing the plus button. This feature is meant for gamers, but it works beautifully for drawing and editing graphics. Aside from a drawing tablet, it's hard to beat this mouse for a drawing tool. You have instant control over the responsiveness of your mouse without taking your hand off the mouse itself.■

Nicholas Petreley is Editor in Chief of *Linux Journal* and a former programmer, teacher, analyst and consultant who has been working with and writing about Linux for more than ten years.

Hacking Cell Phones via Bluetooth Tools under Linux

Ditch the funky USB cable and connect to your cell phone and other devices with Bluetooth.

PATRICK M. DAVILA

Many modern cell phones can take pictures and short video clips. This is great for taking candid photos and sharing them with friends and family at a later time. Modern cell phones also commonly have replaceable themes and ringtones, which are, in essence, 30-second audio clips. On newer phones, ringtones often are MP3 files, replacing older, nonstandard formats. According to many wireless providers, the only way to get your media off the phone is to e-mail yourself using the provider's data transfer service. This can become rather expensive, especially if you take a lot of photos. Some cell phones come with a standard USB connection and can be mounted as USB mass storage devices. Many cell phones, including the phone I currently own (Motorola V551), utilize a nonstandard USB connection. So, instead of purchasing a specialized USB cable, I decided to connect to it using a standard USB Bluetooth dongle. The dongle should work with any phone that supports the Bluetooth protocol.

Bluetooth is a wireless communication protocol that enables transfer speeds of up to 3Mbps and up to 100 meters in distance. It's currently used in various devices, such as printers, keyboards, mice, wireless headsets and GPS receivers. The Bluetooth specification was first developed in 1994 by Ericsson Mobile Platform engineers Sven Mattison and Jaap Haartsen, and it has been adopted as a standard by more than 6,000 companies worldwide. Most modern cell phones now include some sort of Bluetooth wireless capability.

By using Bluetooth, you can push and pull files to and from a cell phone, which is handy if you want to archive photos taken with your cell phone and don't want to pay $.25 or more to e-mail yourself each file. This also allows you to create ringtones from songs that you own in MP3 format (many cell phones let you play standard MP3 audio files). And, some users will be able to connect to the Internet via PPP and other utilities. In this article, we mostly are concerned with the tools available in Linux to access files on your Bluetooth-enabled cell phone. This includes the low-level command-line utilities as well as the GUI applications available under KDE and GNOME. Comparable functionality from a proprietary Windows/Mac application will cost you anywhere from $15–$40 US.

Hardware

Besides having a Bluetooth-enabled cell phone, you need to have a Bluetooth connection for your computer. Many newer laptops come with an internal Bluetooth card, although most desktops do not. If you don't have built-in Bluetooth capability, you need to purchase some sort of Bluetooth transmitter. The most common are USB 1.1 and 2.0 dongles, which are available from retail electronics outlets and mobile phone vendors. I purchased a D-Link DBT-120 Bluetooth dongle for roughly $25 US. Even though Bluetooth is an industry standard, it's probably a good idea to check hardware-compatibility lists to make sure a particular model will work under Linux. The DBT-120 supports speeds of 723Kbps asymmetrical and 433Kbps symmetrical. Although not blazingly fast, these speeds are reasonable for small file transfers. Newer dongle models will give you up to 3Mbps.

In the past, Linux users would have to recompile their kernel to get the Bluetooth modules working correctly. Fortunately, almost all newer distributions of Linux now come with Bluetooth enabled in the kernel by default.

BlueZ Command-Line Tools

The BlueZ package is the default toolset for Bluetooth in Linux. You may need to install the Bluetooth utilities before you get started. In Debian, for example, you should install the bluez-utils package at the very least.

Once you have your Bluetooth dongle connected to your machine, you should check whether all the kernel modules are loading correctly:

```
root@host#  lsmod | grep blue
```

And, you should see something like the following:

```
bluetooth               41060  14 hidp,rfcomm,l2cap,hci_usb
```

Then, type:

```
root@host# dmesg | grep Blue
```

which should produce something like this:

```
Bluetooth: Core ver 2.8
Bluetooth: HCI device and connection manager initialized
Bluetooth: HCI socket layer initialized
Bluetooth: HCI USB driver ver 2.9
Bluetooth: L2CAP ver 2.8
Bluetooth: L2CAP socket layer initialized
Bluetooth: RFCOMM socket layer initialized
Bluetooth: RFCOMM TTY layer initialized
Bluetooth: RFCOMM ver 1.7
Bluetooth: HIDP (Human Interface Emulation) ver 1.1
```

Everything looks good. Next, you'll need some information about your system:

```
root@host# hciconfig -a
```

The above spits out a bunch of lines. You are looking for this line

in particular:

```
Name: 'BlueZ mycomputername'
```

In some cases, you won't see the BlueZ portion, and the computer name may have a trailing -0.

Next, open the file /etc/bluetooth/hcid.conf with your favorite text editor, and change the following line:

```
name "BlueZ %h (%d)";
```

to:

```
name "BlueZ mycomputername";
```

Next, create a PIN for the computer to access the cell phone. Open the /etc/bluetooth/pin file, and add the following:

```
PIN:1234
```

The actual PIN number can be anything you like, and it may not be included in a separate file, depending on your distribution. It may be part of your hcid.conf file and called a passkey instead of a PIN.

The reason for the PIN number is that Bluetooth devices need to be paired or tethered together. This is a standard Bluetooth security measure to prevent unwanted connections between devices. The first time you connect to your phone via Bluetooth, the phone notifies you that a connection is being attempted and prompts you for a PIN number. If the PIN number entered on the phone does not match the PIN in the configuration file, the connection will be rejected. Most cell phones will give you the option to connect devices automatically on subsequent connections or prompt for intervention.

Next, restart the Bluetooth server on the connecting computer:

```
root@host# /etc/init.d/bluetooth restart
```

Some cell phones have a "Find Me" mode that needs to be turned on when scanning. Take a look in your particular phone's Bluetooth connection menu. Now, you're ready to see if the Bluetooth dongle can see your cell phone or any other nearby devices with Bluetooth capability. So, type:

```
root@host# hcitool scan
```

which returns the following:

```
Scanning ...
        00:0F:86:89:EC:3D       Blackberry 7290
        00:14:9A:C9:BB:62       Motorola Phone
        00:16:CB:2A:7D:DB     Mac_1
```

Each device name is listed with its MAC address.

Mac_1 is my USB Bluetooth dongle. Motorola Phone is the name I have designated in the phone's Bluetooth setup menus. BlackBerry 7290 is a nearby device that the my dongle is picking up as well.

Another useful command is `sdptool search DUN`. This provides detailed information for your device:

```
Inquiring ...
Searching for DUN on 00:16:CB:2A:7D:DB ...
Searching for DUN on 00:14:9A:C9:BB:62 ...
Service Name: Dial-up networking Gateway
Service Description: Dial-up networking Gateway
Service Provider: Generic Cellphone Service
Service RecHandle: 0x10001
Service Class ID List:
  "Dialup Networking" (0x1103)
Protocol Descriptor List:
  "L2CAP" (0x0100)
  "RFCOMM" (0x0003)
    Channel: 1
Language Base Attr List:
  code_ISO639: 0x656e
  encoding:    0x6a
  base_offset: 0x100
  code_ISO639: 0x6672
  encoding:    0x6a
  base_offset: 0xd800
  code_ISO639: 0x6573
  encoding:    0x6a
  base_offset: 0xd803
  code_ISO639: 0x7074
  encoding:    0x6a
  base_offset: 0xd806
Profile Descriptor List:
  "Dialup Networking" (0x1103)
    Version: 0x0100
```

Now, you're ready to start pushing and pulling files to/from your cell phone. Let's investigate the available tools in both KDE and GNOME.

KDE Bluetooth Framework

The KDE Bluetooth framework is built on the BlueZ stack and can utilize all the functionality of the command-line tools in an intuitive GUI interface. Originally, it was an add-on application, but because of the proliferation of Bluetooth devices, it has been merged into the baseline KDE desktop. The easiest way to access your data is through the Konqueror file manager. Once the Bluetooth dongle is plugged in to your computer, you should see the Bluetooth icon appear on the Kicker panel. Open Konqueror, and enter `bluetooth:///` in the navigation toolbar.

You should see a listing of nearby devices that are Bluetooth-enabled (Figure 1). Click on your phone, and you should see a listing of available services (Figure 2). For pushing and pulling files, we're most concerned with OBEX File Transfer and OBEX Object Push. Selecting OBEX File Transfer shows the media file folders residing on your device (Figure 3). In my case, I have separate folders for audio, pictures and video clips. Clicking on any of the folders should reveal the files currently on your device. Now you simply can copy/move a file to your home folder on your computer by doing a copy and paste in Konqueror. The first time you do this, you'll be prompted for a PIN number on the cell phone. Subsequent file transfers will not require a PIN unless it has been specified in the cell phone's Bluetooth setup. If you want to copy files from your computer to your device, select OBEX Object Push.

The system displays a pop-up asking if you want to open the

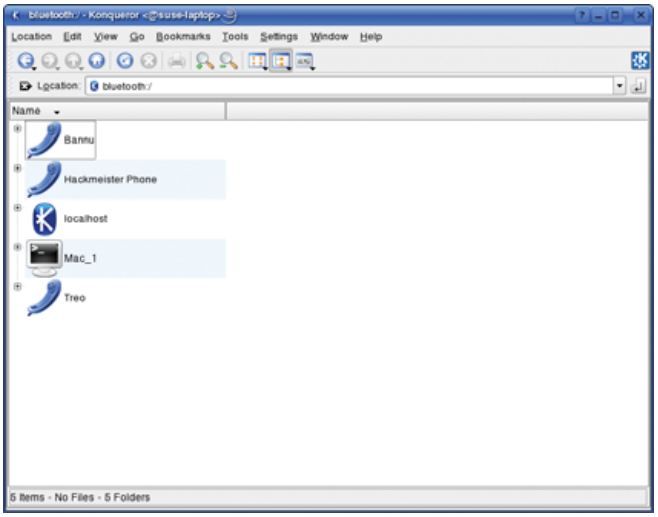

Figure 1. Install KDEBluetooth to make Konqueror aware of Bluetooth devices.

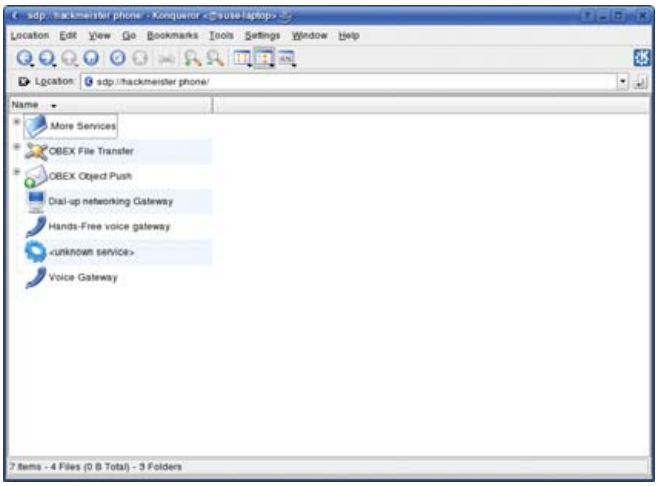

Figure 2. Select your device to get to OBEX Object Push.

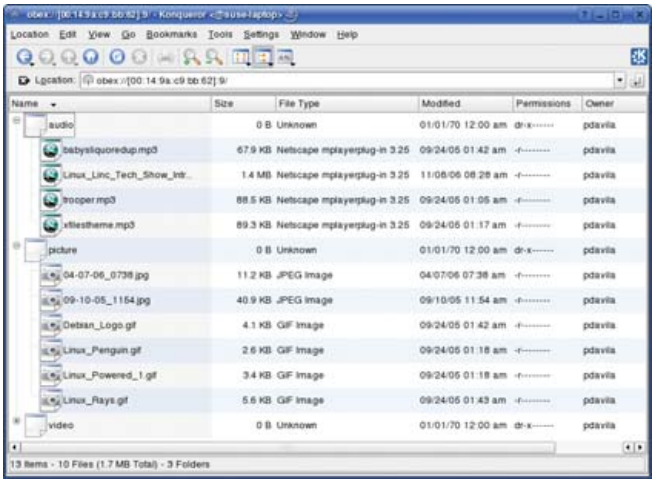

Figure 3. The obex option is where you can launch the OBEX Object Push client.

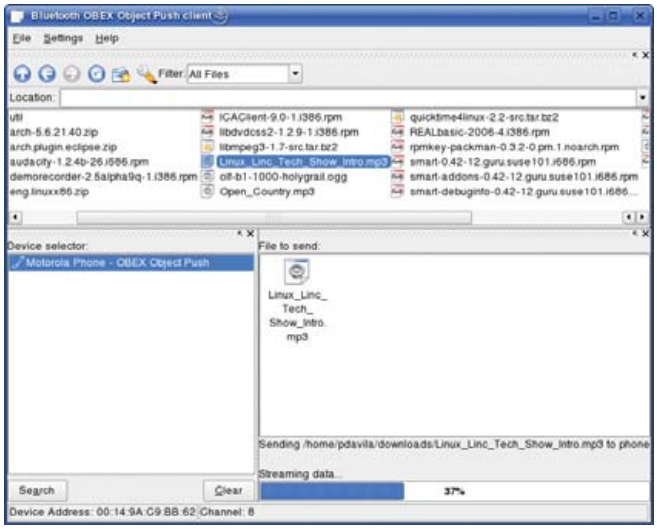

Figure 4. The OBEX Object Push client lets you transfer files with drag and drop.

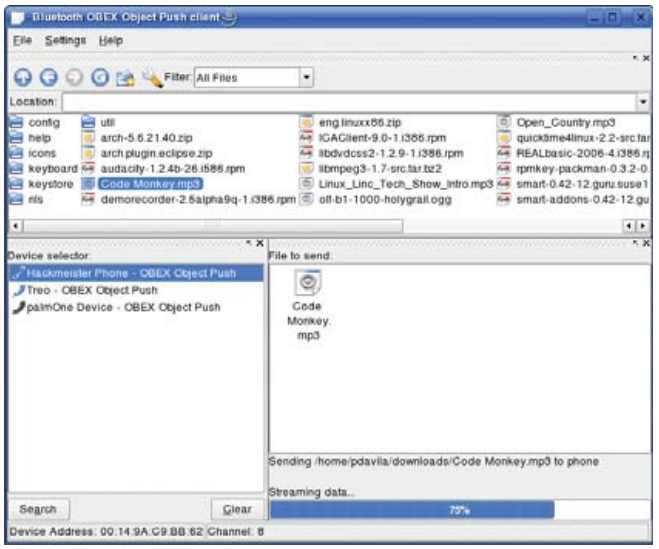

Figure 5. The OBEX Object Push Client with Multiple Devices Available

kbluetooth client. Select yes, and you should see your device list in the left-hand (device selector) column of the client (Figure 4). The top of the client application shows your system folders. The right-hand side of the screen contains a blank area that's titled Files to send. Simply go into one of your system folders and click the file you want to copy to your phone. Now, drag and drop it into the Files to send pane. Click the Send button. Your phone will alert you and ask if you want to accept the file transfer from your computer. Click yes, and the file transfer starts (Figure 5). Once the file transfer is complete, you should be able to find the file on your cell phone or PDA. It doesn't get any easier than that. Konqueror once again proves what a fantastic file manager it is.

KMobileTools

The KMobileTools application is still considered alpha software, but it looks very promising. It's an add-on KDE application that lets you back

GENSTOR
Linux - FreeBSD - x86 Solaris - MS etc.
Simplifying Business
Flexibility
Performance
Investment Protection
Assurance
AMD Opteron™ Processor
The AMD Opteron™ processor with Direct Connect Architecture scales from 1P/2-core up to 8P/16-core across a single industry-standard platform without external logic, allowing for maximum versatility and lowering overall system cost.
Quad CPU
Up to 8 Cores - in 1U
Ideal for high density clustering in standard 1U form factor. Up to 8 Cores for high CPU needs. Comes pre-configured with OS of your choice.
Features:
- 1U rack-optimized chassis (1.75in.)
- Up to 4 Dual Core AMD Opteron™ processors
- Up to 8 Cores Per 1U rackspace
- Up to 64GB DDR2 667 & 533 Mhz SDRAM
- Dual-port Gigabit Ethernet Per Node
- 3 SATA or SCSI Removable HDD Per Node
- 1 (x8) PCI-Express
Servers : : Storage : : Appliances
Genstor Systems, Inc.
780 Montague Expy. # 604
San Jose, CA 95131
www.genstor.com
Email: sales@genstor.com
Ph: 1-877-25 SERVER or 1-408-383-0120
AMD 64 Opteron
AMD, AMD Opteron and the AMD logo are trademarks or registered trademarks of Advanced Micro devices, Inc. Product pictures only for reference.

up, edit and import your cell phone's contact list to your computer using the KAddressbook application. This lets you synchronize your contact list on your cell phone or PDA to your e-mail contact list. You no longer have to maintain separate contact lists for your cell phone and e-mail clients.

Another interesting feature is the ability to control your cell phone remotely from the computer. You can have the cell phone dial a number saved in the contact list or type in another via the dial-out number field. You also can send SMS messages through your cell phone. This is handy if you need to send a long message and don't want to kill your thumbs by typing it with the phone's keypad. KMobileTools also displays the cell phone's battery charge percentage and signal strength status.

GNOME Bluetooth Subsystem

Not to be outdone by KDE, the GNOME desktop also provides a useful set of tools for file transfers to your Bluetooth devices. There is a nice GNOME-based front end to hcitool that can be started by typing `gnome-bluetooth-admin` from a terminal session. Click the Scan button, and you should see a listing of nearby Bluetooth devices. Open the Nautilus file manger, and type `bluetooth:///` in the navigation toolbar. An icon representing your phone or PDA should appear. Next, open another Nautilus session, drag the icons of the files you want to transfer and drop them onto the Bluetooth device's icon. You should see a message on your cell phone asking if you want to accept the files from your computer. You also can right-click on the icon for the file you want to send and select Send via Bluetooth to initialize the file transfer. Once the transfer is complete, you should see the new file on your cell phone.

Another great application is Phone Manger, which is GNOME's answer to KMobileTools. Phone Manager lets you monitor your device's status, send and read SMS messages to and from your phone and integrate your contact list with Evolution's address book.

GET REAL – at – ISPCON FALL 2007

October 16-18, 2007
San Jose McEnery Convention Center
San Jose, CA • www.ispcon.com

ISPCON is where the service provider industry goes to GET REAL about the future of their businesses. Isn't it time you increase ARPU? Pump up profits? Optimize your operations? Find out what's next?

Whatever your plans, ISPCON will help you GET REAL about them. This is the industry's only forum where peers learn from peers in real one-on-one discussion about what works, what doesn't and what's next.

GET REAL! Get to ISPCON!

Register online and use Customer Code LNXF7 to get your FREE EXHIBIT HALL PASS!

Creating Ringtones with Audacity

Now that you can push and pull files to and from your cell phone, let's create some ringtones. Why create a custom ringtone? You can be lazy and buy them for $2.50 US each from your service provider, or you can have some fun, learn something new and save yourself some money. Personally, I prefer the latter. Many newer cell phones let you play standard MP3 files without any tweaking whatsoever.

If your phone has limited storage space, it might be prudent to reduce the size of the audio file in Audacity. The first step is to import your song into a new project in Audacity. Trim the song to an approximately 30-second clip of a section of the song that you like. Next, modify the clip so it's in mono. (This shrinks the file size further.) Click the track's filename, and select Split Stereo Track. Delete the right channel of the clip by clicking the X in the upper-left corner of the track. Select the remaining channel, and click on the down arrow next to the track name. You will see a drop-down menu containing entries for Mono, Left Channel and Right Channel. Select Mono.

Next, you need to normalize the clip. First, press Ctrl-A to select the whole track. Then, select the normalize entry in the Effect drop-down menu in Audacity's main toolbar.

Now, you're ready to encode the MP3 file of your ringtone. Click File, and select Export Selection as MP3. Take the new MP3 file and simply move it into the audio folder on your cell phone as described earlier in the article. This is a great method for creating customized ringtones using a FOSS application.

Using Your Cell Phone as a Modem

It is possible to use your Bluetooth-enabled cell phone or PDA as a modem if it has GPRS or EDGE network capability. One caveat is that many cell-phone providers lock this functionality out of their phones unless the customer has a data plan with them. Data plan prices vary significantly, depending upon how much throughput you are allowed and which carrier you use. I obtained an introductory plan for $25 US per month. An unlimited data plan probably costs in the range of $60 US a month. This is a good option for road warriors who might be traveling into areas where broadband is not readily available, but cell-phone service is. Nowhere on my cell-phone provider's Web site did it mention that this service is available to Linux users. However, the provider offers a custom application for Windows users. The application provides users with Internet access using their service. Mac OS X users were relegated to a How-To document on the customer support Web site.

It was only when I requested to talk to the technical support that the provider openly acknowledged that connectivity under Linux was possible. I was then provided with a custom script that calls the WvDial PPP dial-up application. The script contained specific initialization parameters to pass to the phone and connect to the dial-up service. The technical support person at first claimed I would be able to use this service only if I connected the cell phone to my laptop via a custom Motorola USB cable. Conveniently, they sell this cable for $30

US and promptly offered to sell it to me. I communicated that I was able to pass data to/from my phone seamlessly using a generic Bluetooth dongle. At this point, the technical support person admitted that the USB cable was not necessary. I then paid for a month's worth of service and was informed that customer support does not offer any assistance to Linux users. Nothing new there. Subsequently, I sent an e-mail to customer support asking that they acknowledge on their support pages that the data service works flawlessly with Linux. I have not heard back from them at the time of this writing.

Conclusion

The tools available for Bluetooth devices under Linux are fairly robust. They provide all of the functionality of similarly closed applications available for Windows and Mac at no cost. Once again, the Linux community should be commended for developing such great programs. The command-line tools provide a good insight into the protocols and services used in accessing Bluetooth devices. The GUI applications are intuitive enough that average Linux users can put them to good use in short order. Linux is right in the mix in the growing field of communication devices and the applications that manage them.■

Patrick M. Davila is a programmer and has been a Linux enthusiast since 1998. He is a co-host of *The Linux Link Tech Show*, the longest continually running Linux show on the Internet. When not tinkering with Linux, Patrick enjoys music, beer, soccer and spending time with his family.

Resources

BlueZ Bluetooth Protocol Stack for Linux: **www.bluez.org**

KDE Bluetooth Framework: **bluetooth.kmobiletools.org**

GNOME Bluetooth Subsystem How-To: **usefulinc.com/software/gnome-bluetooth**

Gentoo Bluetooth Wiki: **gentoo-wiki.com/HOWTO_mobile_phone,_Bluetooth_and_GNOME**

General Linux and Bluetooth Links: **www.holtmann.org/linux/bluetooth**

KMobileTools Application: **www.kmobiletools.org**

GNOME Phone Manager: **live.gnome.org/PhoneManager**

Audacity Audio Editor: **audacity.sourceforge.net**

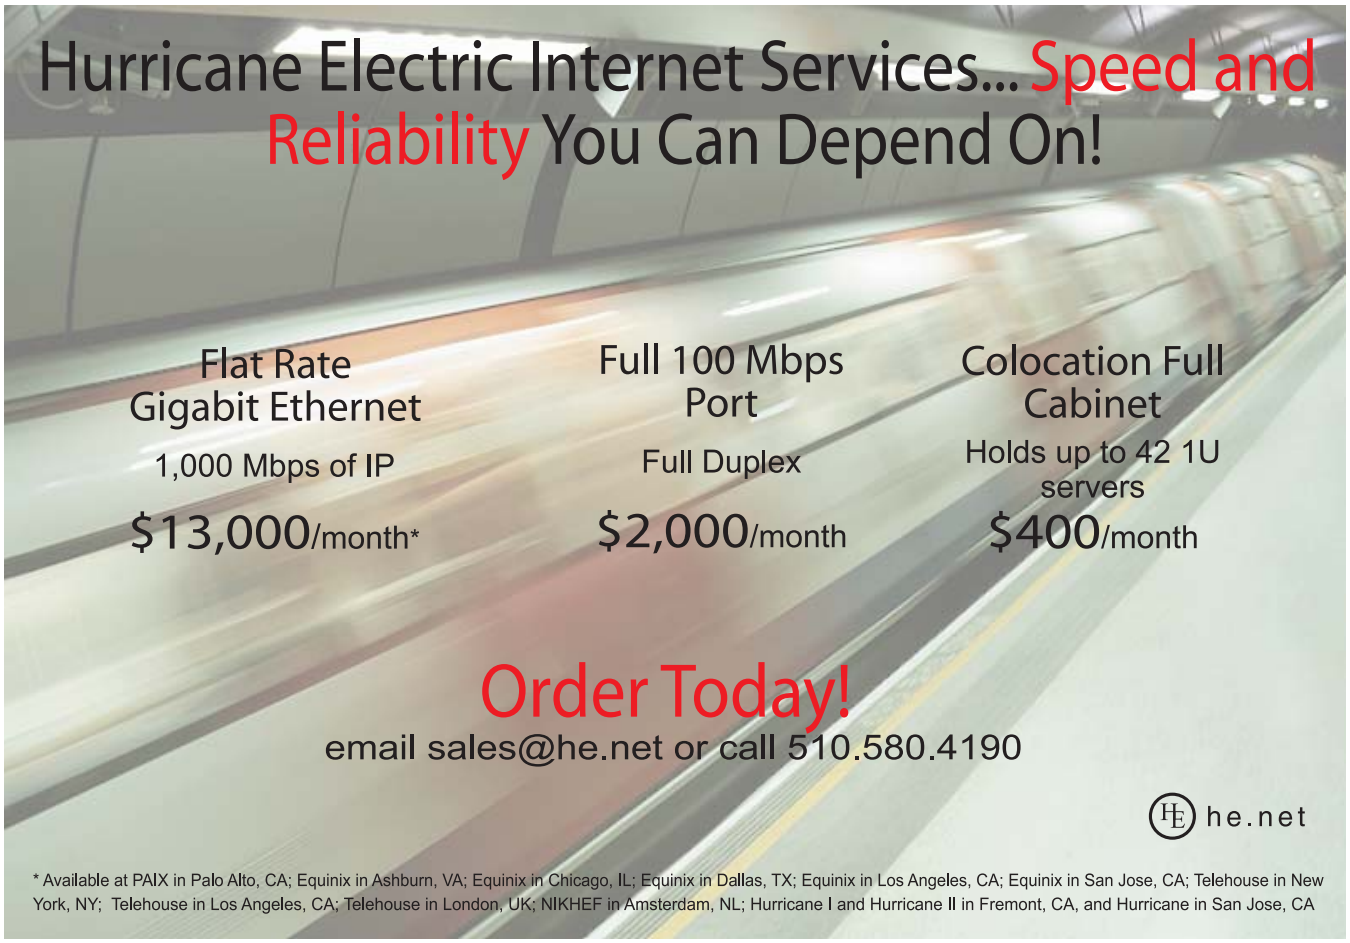

Fedora Directory Server: the Evolution of Linux Authentication

Check out Fedora Directory Server to authenticate your clients without licensing fees.

JERAMIAH BOWLING

In most enterprise networks today, centralized authentication is a basic security paradigm. In the Linux realm, OpenLDAP has been king of the hill for many years, but for those unfamiliar with the LDAP command-line interface (CLI), it can be a painstaking process to deploy. Enter the Fedora Directory Server (FDS). Released under the GPL in June 2005 as Fedora Directory Server 7.1 (changed to version 1.0 in December of the same year), FDS has roots in both the Netscape Directory Server Project and its sister product, the Red Hat Directory Server (RHDS). Some of FDS's notable features are its easy-to-use Java-based Administration Console, support for LDAP3 and Active Directory integration. By far, the most attractive feature of FDS is Multi-Master Replication (MMR). MMR allows for multiple servers to maintain the same directory information, so that the loss of one server does not bring the directory down as it would in a master-slave environment.

Getting an FDS server up and running has its ups and downs. Once the server is operational, however, Red Hat makes it easy to administer your directory and connect native Fedora clients. In addition to providing network authentication, you easily can extend FDS functionality across other applications, such as NFS, Samba, Sendmail, Postfix and others. In this article, we focus solely on using FDS for network authentication and implementing MMR.

Installation

To begin, download a Fedora 6 ISO readily available from one of the many Fedora mirrors. FDS has low hardware requirements—500MHz with 256MB RAM and 3GB or more space. I recommend at least a 1GHz processor or above with 512MB or more memory and 20GB or more of disk space. This configuration should perform well enough to support small businesses up to enterprises with thousands users. As for supported operating systems, not surprisingly, Red Hat lists Fedora and Red Hat flavors of Linux. HP-UX and Solaris also are supported. With your bootable ISO CD, start the Fedora 6 installation process, and select your desired system preferences and packages. Make sure to select Apache during installation. Set your host and DNS information during the install, using a fully qualified domain name (FQDN). You also can set this information post-install, but it is critical that your host information is configured properly. If you plan to use a firewall, you need to enter two ports to allow LDAP (389 default) and the Admin Console (default is random port). For the servers used here, I chose ports 3891 and 3892 because of an existing LDAP installation in my environment. Fedora also natively supports Security-Enhanced Linux (SELinux), a policy-based lock-down tool, if you choose to use it. If you want to use SELinux, you must choose the Permissive Policy.

Once your Fedora 6 server is up, download and install the latest RPM of Fedora Directory Server from the FDS site (it is not included in the Fedora 6 distribution). Running the RPM unpacks the program files to /opt/fedora-ds. At this point, download and install the current Java Runtime Environment (JRE) .bin file from Sun before running the local setup of FDS. To keep files in the same place, I created an /opt/java directory and downloaded and ran the .bin file from there. After Java is installed, replace the existing soft link to Java in /etc/alternatives with a link to your new Java installation. The following syntax does this:

```
cd /etc/alternatives
rm ./java
ln -sf /opt/java/jre1.5.0_09/bin/java java
```

Next, configure Apache to start on boot with the chkconfig command:

```
chkconfig -level 345 httpd on
```

Then, start the service by typing:

```
service httpd start
```

Now, with the useradd command, create an account named fedorauser under which FDS will run. After creating the account, run `/opt/fedora-ds/setup/setup` to launch the FDS installation script. Before continuing, you may be presented with several error messages about non-optimal configuration issues, but in most cases, you can answer yes to get past them and start the setup process. Once started, select the default Install Mode 2 - Typical. Accept all defaults during installation except for the Server and Group IDs, for which we are using the fedorauser account (Figure 1), and if you want to customize the ports as we have here, set those to the correct numbers (3891/3892). You also may want to use the same passwords for both the configuration Admin and Directory Manager accounts created during setup.

When setup is complete, use the syntax returned from the script to access the admin console (`./startconsole -u admin http://fullyqualifieddomainname:port`) using the Administrator account (default name is Admin) you specified during the FDS setup.

INTEROP®

NEW YORK | OCTOBER 22–26, 2007

BUSINESS. TECHNOLOGY.
ONE WEEK. ONE PLACE.

- APPLICATION DELIVERY
- BRANCH OFFICE
- COLLABORATION
- ENTERPRISE 2.0
- FIXED MOBILE CONVERGENCE
- GREEN IT
- NETWORKING & SERVICES
- OPTIMIZATION
- SECURITY
- SOA
- STORAGE
- UNIFIED COMMUNICATIONS
- VIRTUALIZATION
- VOIP
- WIRELESS & MOBILITY

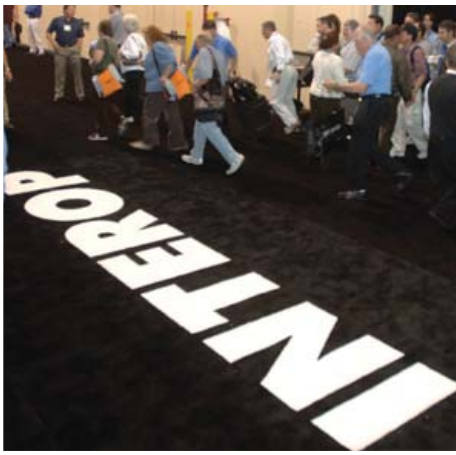

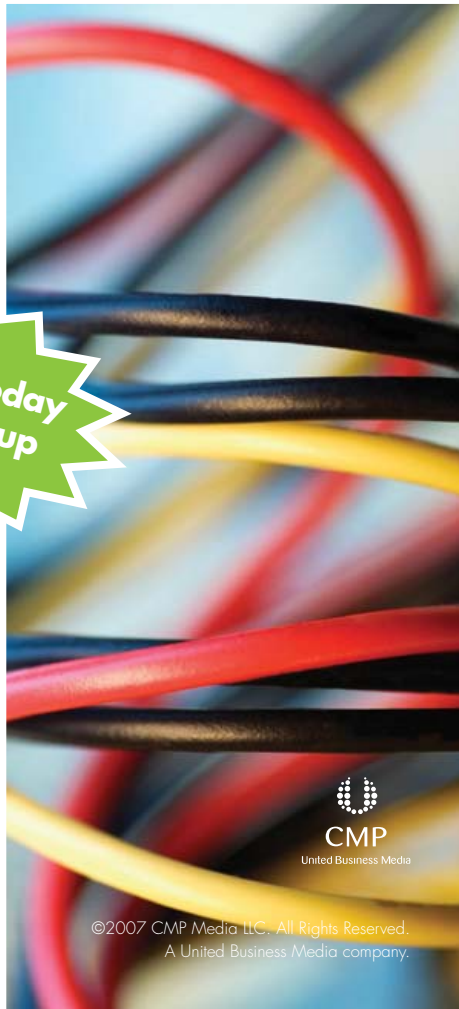

Register Today and Save up to $500

DON'T MISS THE LEADING BUSINESS TECHNOLOGY EVENT

This Fall Interop will bring IT and business leaders together in New York to see the latest technologies in action. Whether you need to evaluate a product, solve a challenge or just stay ahead of the curve, Interop will help you move your business forward.

Register Today.

Register with priority code MLSUNN26 to save up to $500 off any conference pass at www.interop.com.

CMP
United Business Media

©2007 CMP Media LLC. All Rights Reserved.
A United Business Media company.

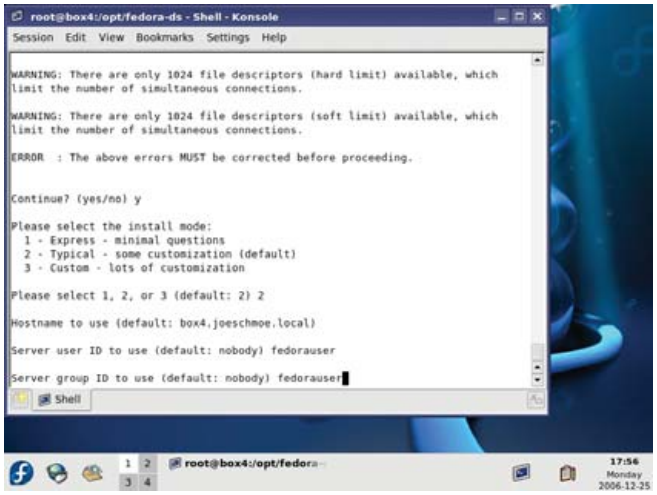

Figure 1. Install Options

You always can call the Admin console using this same syntax from /opt/fedora-ds.

At this point, you have a functioning directory server. The final step is to create a startup script or directly edit /etc/rc.d/rc.local to start the slapd process and the admin server when the machine starts. Here is an example of a script that does this:

```
/opt/fedora-ds/slapd-yourservername/start-slapd
/opt/fedora-ds/start-admin &
```

Directory Structure and Management

Looking at the Administration console, you will see server information on the Default tab (Figure 2) and a search utility on the second tab. On the Servers and Applications tab, expand your server name to display the two server consoles that are accessible: the Administration Server and the Directory Server. Double-click the Directory Server icon, and you are taken to the Tasks tab of the Directory Server (Figure 3). From here, you can start and stop directory services, manage certifi-

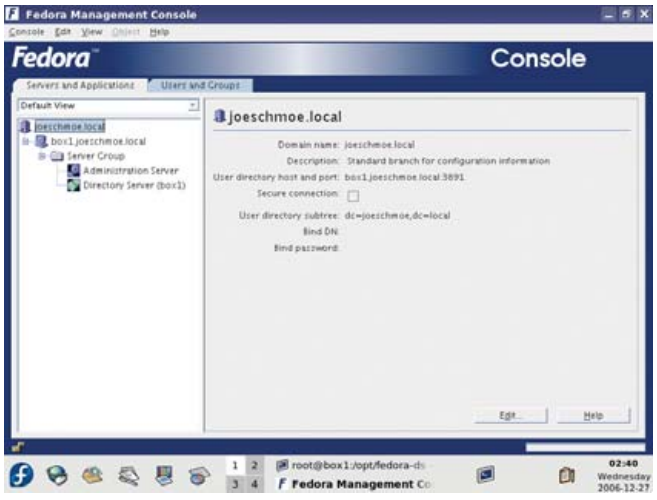

Figure 2. Main Console

Figure 3. Tasks Tab

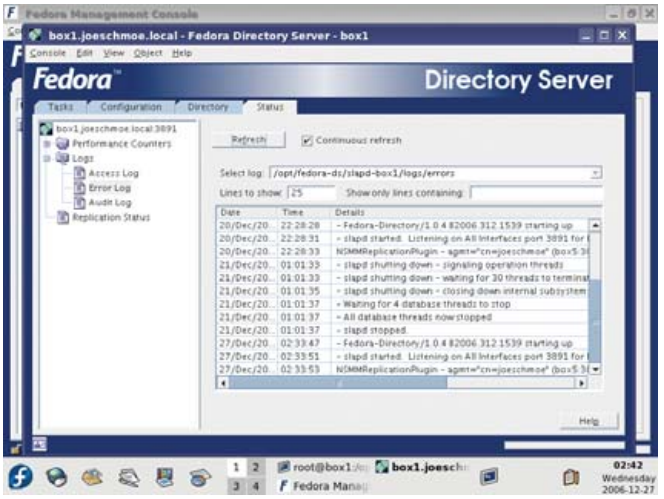

Figure 4. Main Logs

cates, import/export directory databases and back up your server. The backup feature is one of the highlights of the FDS console. It lets you back up any directory database on your server without any knowledge of the CLI. The only caveat is that you still need to use the CLI to schedule backups.

On the Status tab, you can view the appropriate logs to monitor activity or diagnose problems with FDS (Figure 4). Simply expand one of the logs in the left pane to display its output in the right pane. Use the Continuous Refresh option to view logs in real time.

From the Configuration tab, you can define replicas (copies of the directory) and synchronization options with other servers, edit schema, initialize databases and define options for logs and plugins. The Directory tab is laid out by the domains for which the server hosts information. The Netscape folder is created automatically by the installation and stores information about your directory. The config folder is used by FDS to store information about the server's operation. In most cases, it should be left alone, but we will use it when we set up replication.

LinuxWorld 07

CONFERENCE & EXPO

23-24 October 2007 • Olympia 2 • London

Open for business

LinuxWorld Conference & Expo2007 is the premier event focusing exclusively on Open Source and Linux technologies. The show, which attracts technical decision makers and business professionals will showcase the latest technology, debate the use of Open Source in the desktop environment, touch on the use of Linux in the datacentre and discuss the development of a common Linux-based platform for mobile phones.

Conference

Attend the groundbreaking conference, where you can see world leading technical experts discuss the industry's more relevant topics. Aimed at technical and business professionals, the 2007 conference will offers:

- 2 full days of content, covering many fields
- 2 conference/masterclass streams
- Indepth technical sessions
- Business focused seminars

Book your place now by registering at www.linuxworldexpo.co.uk

Expo

As well as being able to visit over 60 exhibitors, showcasing new products and services, the Expo offers wide reaching benefits to you, the visitor:

- On the floor business briefings, showcasing real life examples of advantages, applications and integration of Linux and Open Source in the business
- Great Linux Debate
- Take LPI Examinations at a special visitor discount
- Meet developers and technical experts in the .org Village

Register now for free entry at www.linuxworldexpo.co.uk

IDG WORLD EXPO

2007 IDG Word Expo Corp. All rights reserved. LinuxWorld and LinuxWorld Conference & Expo are registered trademarks of International Data Group Inc. All other trademarks contained herein are the property of the respective owners. All trademarks acknowledged. E&OE. Programme may be subject to change. Correct at time of press.

Before creating your directory structure, you should carefully consider how you want to organize your users in the directory. There is not enough space here for a decent primer on directory planning, but I typically use Organizational Units (OUs) to segment people grouped together by common security needs, such as by department (for example, Accounting). Then, within an OU, I create users and groups for simplifying permissions or creating e-mail address lists (for example, Account Managers). FDS also supports roles, which essentially are templates for new users. This strategy is not hard and fast, and you certainly can use the default domain directory structure created during installation for most of your needs (Figure 5). Whatever strategy you choose, you should consult the Red Hat documentation prior to deployment.

To create a new user, right-click an OU under your domain and click on New→User to bring up the New User screen (Figure 5). Fill in the appropriate entries. At minimum, you need a First Name, Last Name and Common Name (cn), which is created from your first and last name. Your login ID is created automatically from your first initial and your last name. You also can enter an e-mail address and a password. From the options in the left pane of the User window, you can add NT User information, Posix Account information and activate/de-activate the account. You can implement a Password Policy from the Directory tab by right-clicking a domain or OU and selecting Manage Password Policy. However, I could not get this feature to work properly.

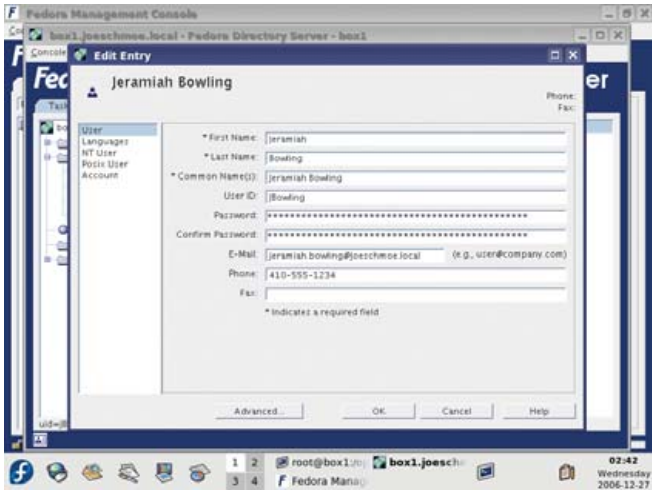

Figure 5. User Screen

Replication

Setting up replication in FDS is a relatively painless process. In our scenario, we configure two-way multi-master replication, but Red Hat supports up to four-way. Because we already have one server (server one) in operation, we need another system (server two) configured the same way (Fedora, FDS) with which to replicate. Use the same settings used on server one (other than hostname) to install and configure Fedora 6/FDS. With both servers up, verify time synchronization and name resolution between them. The servers must have synced time or be relatively close in their offset, and they must be able to resolve the other's hostname for replication to work. You can use IP addresses, configure /etc/hosts or use DNS. I recommend having DNS and NTP in place already to make life easier.

The next step is creating a Replication Manager account to bind one server's directory to the other and vice versa. On the Directory tab of the Directory Server console, create a user account under the config folder (off the main root, not your domain), and give it the name Replication Manager, which should create a login ID (uid) of RManager. Use the same password for the RManager on both servers/directories. On server one, click the Configuration tab and then the Replication folder. In the right pane, click Enable Changelog. Click the Use Default button to fill in the default path under your slapd-servername directory. Click Save. Next, expand the Replication folder and click on the userroot database. In the right pane, click on enable replica, and select Multiple Master under the Replica Role section (Figure 6). Enter a unique Replica ID. This number must be different on each server involved in replication. Scroll down to the update section below, and in the Enter a new supplier DN: textbox, enter the following:

```
uid=RManager,cn=config
```

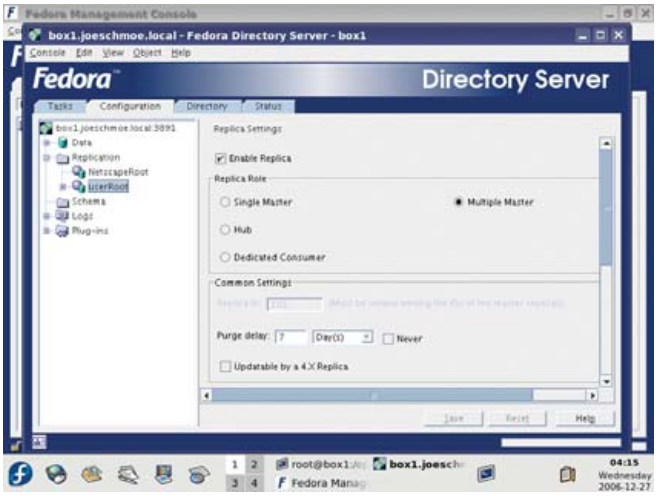

Figure 6. Setting Up Replica

Click Add, and then the Save button at the bottom. On server two, perform the same steps as just completed for creating the RManager account in the config folder, enabling changelogs and creating a Multiple Master Replica (with a different Replica ID).

Now, you need to set up a replication agreement back on server one. From the Configuration tab of the Directory Server, right-click the userroot database, and select New Replication Agreement. A wizard then guides you through the process. Enter a name and description for the agreement (Figure 7). On the Source and Destination screen under Consumer, click the Other button to add server two as a consumer. You must use the FQDN and the correct port (3891 in our case). Use Simple Authentication in the Connection box, and enter the full context (uid=Rmanager,cn=config) of the RManager account and the password for the account (Figure 8). On the next screen, check off the box to enable Fractional Replication to send only deltas rather than the entire directory database when changes occur, which is very useful over WAN connections. On the Schedule screen, select Always keep directories in sync to keep. You also could choose to schedule your updates. On the Initialize Consumer screen, use the Initialize Now option. If you experience any difficulty in initializing a consumer (server two), you can use the Create consumer initialization file option and

copy the resulting ldif folder to server two and import and initialize it from the Directory Server Tasks pane. When you click next, you will see the replication taking place. A summary appears at the end of the process. To verify replication took place, check the Replication and Error logs on the first server for success messages. To complete MMR, set up a replication agreement on the server, listing server one as the consumer, but do not initialize at the end of the wizard. Doing so would overwrite the directory on server one with the directory on server two. With our setup complete, we now have a redundant authentication infrastructure in place, and if we choose, we can add another two Read-Write replicas/Masters.

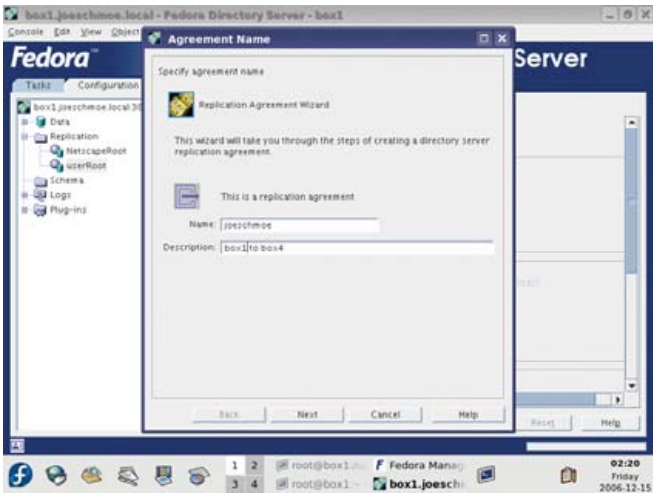

Figure 7. Enter a name and password.

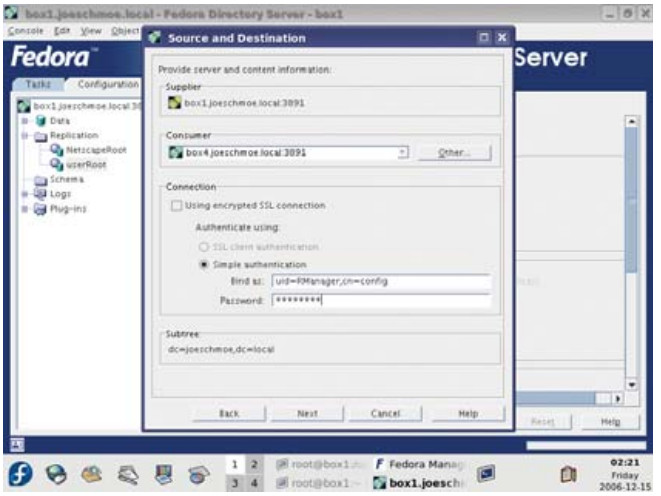

Figure 8. Enter information for the RManager account.

Authenticating Desktop Clients

With our infrastructure in place, we can connect our desktop clients. For our configuration, we use native Fedora 6 clients and Windows XP clients to simulate a mixed environment. Other Linux flavors can connect to FDS, but for space constraints, we won't delve into connecting them. It should be noted that most distributions like Fedora use PAM,

Unbeatable
Price, Quality, & Service

$1995

1U Supermicro 6015B-TV: (Customized System)
2x Xeon Quad-Core E5335 2GHz CPU,
4GB 667mhz FB-DIMM, 2x 250GB HD, Slim DVD...

$2395

2U Supermicro 6025B-TR+B: (Like above spec.)
2x Xeon E5335 CPU, 4GB RAM, 2x 250GB HD
w/ Redundant Power, 8 hotswap bays ....

*Free gadget w/ purchase!

Windows redhat SUSE fedora solaris RoHS COMPLIANT

Customizable system solutions since 1989

Tel: 1-800-875-8590 Fax:408-736-4151

KSC
KING STAR COMPUTER
1259 Reamwood Ave
Sunnyvale, CA 94089
www.kingstarusa.com
Email: sales@kingstarusa.com

intel Xeon inside™

Intel©, Intel© Xeon™, Intel Inside©, Intel© Itanium© and the Intel Inside© logo are trademarks or registered trademarks of Intel Corporation of its subsidiaries in the United States and other countries. Prices and availability subject to change without notice. Not responsible for typographical errors.

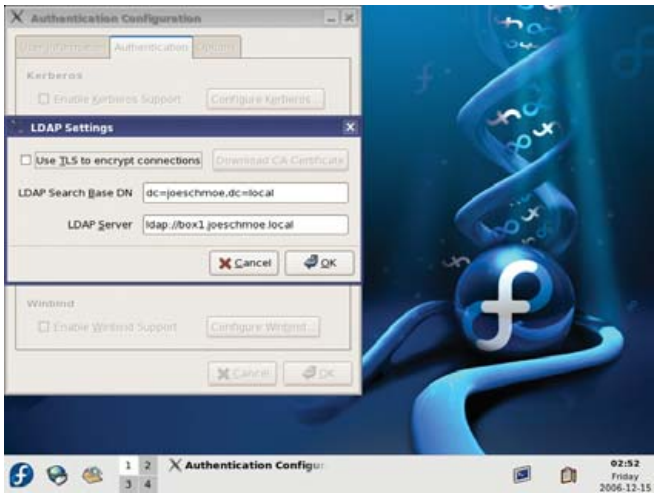

Figure 9. The Built-in Authentication Utility

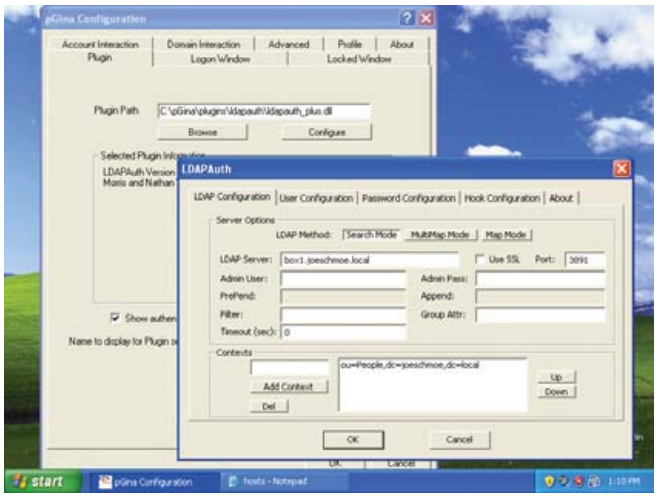

Figure 10. Plugin Tester Tool

the /etc/nsswitch.conf and /etc/ldap.conf files to set up LDAP authentication. Regardless of client type, the user account attempting to log in must contain Posix information in the directory in order to authenticate to the FDS server. To connect Fedora clients, use the built-in Authentication utility available in both GNOME and KDE (Figure 9). The nice thing about the utility is that it does all the work for you. You do not have to edit any of the other files previously mentioned manually. Open the utility and enable LDAP on the User Information tab and the Authentication tabs. Once you click OK to these settings, Fedora updates your nsswitch.conf and /etc/pam.d/system-auth files immediately. Upon reboot, your system now uses PAM instead of your local passwd and shadow files to authenticate users.

During login, the local system pulls the LDAP account's Posix information from FDS and sets the system to match the preferences set on the account regarding home directory and shell options. With a little manual work, you also can use automount locally to authenticate and mount network volumes at login time automatically.

Connecting XP clients is almost as easy. Typically, NT/2000/XP users are forced to use the built-in MSGINA.dll to authenticate to Microsoft networks only. In the past, vendors such as Novell have used their own proprietary clients to work around this, but now the open-source pgina client has solved this problem. To connect 2000/XP clients, download the main pgina zipfile from the project page on SourceForge, and extract the files. For this article, I used version 1.8.4 as I ran into some dll issues with version 1.8.8. You also need to download and extract the Plugin bundle. Run the x86 installer from the extracted files, accepting all default options, but do not start the Configuration Tool at the end. Next, install the LDAPAuth plugin from the extracted Plugin bundle. When done installing, open

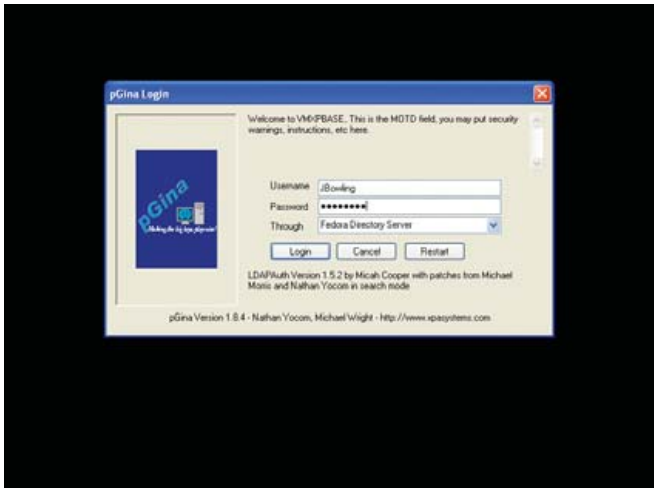

Figure 11. Login Prompt

the Configuration Tool under the Pgina Program Group under the Start menu. On the Plugin tab, browse to your ldapauth_plus.dll in the directory specified during the install. Check off the option to Show authentication method selection box. This gives you the option of logging locally if you run into problems. Without this, the only way to bypass the pgina client is through Safe Mode. Now, click on the Configure button, and enter the LDAP server name, port and context you want pgina to use to search for clients. I suggest using the Search Mode as your LDAP method as it will search the entire directory if it cannot find your user ID. Click OK twice to save your settings. Use the Plugin Tester tool before rebooting to load your client and test connectivity (Figure 10). On the next login, the user will receive the prompt shown in Figure 11.

Conclusions

FDS is a powerful platform, and this article has barely scratched the surface. There simply is not room to squeeze all of FDS's other features, such as encryption or AD synchronization, into a single article. If you are interested in these items or want to know how to extend FDS to other applications, check out the wiki and the how-tos on the project's documentation page for further information. Judging from our simple configuration here, FDS seems evolutionary, not revolutionary. It does not change the way in which LDAP operates at a fundamental level. What it does do is take the complex task of administering LDAP and makes it easier while extending normally commercial features, such as MMR, to open source. By adding pgina into the mix, you can tap further into FDS's flexibility and cost savings without needing to deploy an array of services to connect Windows and Linux clients. So, if you are looking for a simple, reliable and cost-saving alternative to other LDAP products, consider FDS.■

Resources

Main Fedora Site: **fedora.redhat.com**

Fedora ISOs: **fedora.redhat.com/Download**

Fedora Directory Server Site: **directory.fedora.redhat.com**

Main pgina Site: **www.pgina.org**

pgina Downloads:
sourceforge.net/project/showfiles.php?group_id=53525

Jeramiah Bowling has been a systems administrator and network engineer for more than ten years. He works for a regional accounting and auditing firm in Hunt Valley, Maryland, and holds numerous industry certifications including the CISSP. Your comments are welcome at jb50c@yahoo.com.

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.
540 Dado Street, San Jose, CA 95131
www.ironsystems.com

CALL: 1-800-921-IRON

A $7,000 Server Comparison

Tired of x86? See what Linux on Itanium, Sun T1 or POWER5 can do!

PETER ARREMANN

The story of Linux on non-x86 architectures started in 1994 with a port to the now-abandoned Alpha architecture. Other ports quickly followed, and over the years, Linux has gained support for most desktop and server CPU designs. Today, however, only five CPU architectures are promoted actively by their manufacturers as Linux-compatible. This article explores how entry-level servers based non-x86 designs compare to the current x86 systems in the same price range.

CPU Architectures

Comparing the x86 server market is usually fairly boring. The market is split into two camps around the AMD Opteron and the Intel Xeon. The differences between the various server models inside each camp are fairly small. Number of expansion slots, disk count and the features of the remote management solution seem to be the only distinctions. Performance and memory capabilities are determined by the CPU and chipset.

Outside the x86 market, the picture changes. To compete with the established x86 solutions and the massive budget Intel can invest into CPU development, IBM, Sun and the Intel Itanium team have to be innovative and take ideas to new heights.

King of the Hill—x86

The first member of the x86 architecture was the 16-bit 8086 designed by Intel in 1978. Since then, x86 has come a long way. It was extended to 32-bit with the i386 and more recently to 64-bit with the AMD64/EMT64. Despite these extensions, all x86 designs have remained backward-compatible, and even the newest quad-core Xeons and Opterons still run DOS.

This backward compatibility has allowed the x86 processors to become the standard for desktops and also to dominate the market for smaller servers. It is, however, also the reason for much of the criticism that Intel and AMD receive.

In 1978, ideas like pipelining, out-of-order execution and branch prediction were known but did not influence the design of the x86 instruction set. Today, these features are part of most CPUs, and a lot of effort is required to implement these features. This increases complexity, and in many cases, optimal performance is not possible.

The EPIC Story of Itanium

EPIC (Explicitly Parallel Instruction Computing) is the instruction set used in the Intel Itanium processors. EPIC was codeveloped by HP and Intel as the successor to both the HP PA-RISC line and the Intel x86 processors. The development started in 1994, but after delays and missed performance targets, the project's goals have changed dramatically. Although HP has discontinued the PA-RISC and Alpha architectures and is now selling a full range of Itanium-based servers, Intel continued the development of x86-based processors and now positions the Itanium processor only for high-end applications.

The main idea behind EPIC is that the compiler has a much better understanding of the program code than the CPU does. This additional knowledge about the program can be used to optimize the code at compile time rather than during execution. The reduced need for hardware-based optimization results in simpler architecture. However, the decision also requires more effort from compiler designers and leads to some interesting behavior (see The Compiler Issue sidebar).

CMT—Sun T1

CMT, short for Chip Multi-Threading, is only one of the names describing methods for increasing CPU resource utilization. Instead of relying on larger caches or higher clock speed, CMT increases performance by offering multiple execution threads on a single processor.

CMT can be implemented in two variants. The first method is the use of multiple identical cores that are combined in the same physical package. This allows server manufacturers to deliver more processing power per socket and is implemented in all current architectures.

The second type of CMT is allowing one CPU core to execute multiple threads to increase resource utilization. This can be done by providing dedicated resources to each thread or simply by allowing the primary thread full access and limiting the secondary thread to the resources not used by the primary thread. Intel has implemented this feature in many Pentium 4 CPUs under the brand name of HyperThreading. HyperThreading can speed up execution by up to 20%, but workloads that rely heavily on cache sizes (such as the bzip2 compression discussed later in the article) suffer from having HyperThreading enabled.

The T1 processor that Sun is utilizing in the CoolThreads T1000 and T2000 systems uses both CMT concepts. It has eight cores, and each core is capable of executing four simultaneous threads. To combine such a high number of cores on one chip, Sun has chosen to

The Compiler Issue

GCC is the standard compiler for Linux and many other platforms. However, GCC has a long history of being criticized for lack of optimization for non-x86 platforms. This seems to be especially true for the Itanium platform, as EPIC is the newest instruction set and GCC developers had the least amount of time to optimize the compiler. A whitepaper on Intel's Web site describes about a 25% performance gain when simply translating MySQL with the Intel Compiler vs GCC 4.1.

To verify this claim, we recompiled bzip2 and PostgreSQL 7.4.16 on the HP rx2660. The performance gains were impressive—29% for bzip2 and 21% for PostgreSQL. Hopefully, Intel and HP will continue working with the GCC team on improving performance, because adoption of a closed-source compiler by Red Hat and others is unlikely.

implement very basic cores running at a fairly low clock frequency of 1–1.4GHz. This results in low single-thread execution speed, but Sun is betting on the 32 execution thread to make up for this disadvantage.

Balanced: IBM POWER5

The Power architecture is the big brother of the PowerPC chips used in the current generation of gaming consoles, many embedded systems and, until recently, in Macs. The POWER5 processor supports all PowerPC features and adds a special hypervisor mode. This mode is similar to the new Intel-VT and AMD-Pacifica visualization technologies and allows multiple operating systems to run on the same system.

The POWER5 team at IBM has decided to balance single-core performance with a multicore and multithreading implementation. The result is the POWER5 Quad-Core Module (QCM) used in the 510Q. It has four processing cores and the capability of running two independent threads per core.

In addition to balancing the design, IBM invested heavily into manufacturing technology and automated design tools. This allows IBM to reach high clock speeds and produce top-performing processors with much less effort than its competitors.

The Servers

Reviewers often select servers based on the number of CPUs and memory, and then compare the prices. This works well for an x86-based comparison, but the servers covered in this article are too different to be compared by CPU count or number of memory slots. Instead, this article evaluates the servers based on cost. In other words, what kind of features and performance can $7,000 buy?

All servers were purchased with standard one-year warranty and no operating system. The internal disks are used only for the OS installation. The database and application files are located on an external SCSI disk array connected via an LSI Ultra-320 controller.

Sun Fire T1000

The Sun Fire T1000 is the smallest of the four CoolThreads servers currently sold by Sun. It is a 1U unit and comes with a 1GHz T1 processor. Depending on the configuration, either six or all eight cores are enabled. Eight slots of registered DDR2 memory support configurations from 2 to 32GB.

Four gigabit Ethernet ports and a remote management card called ALOM (Advanced Lights Out Manager) are standard. The ALOM is one of the most easy-to-use and capable remote management methods found on UNIX servers. One PCI-Express slot is available for expansion.

Like most 1U servers, the T1000 has only a single power supply. A single 3.5" SATA drive comes standard. A cold-swap drive tray for two 2.5" disks is available as an option. Hot-swap disks are not available.

The server selected for the review was equipped with eight 1GHz cores, 8GB of RAM and a single 160GB disk. Quoted at $7,322, this configuration was just barely over the target price for this review.

Because the T1 is a complete SPARC V9 implementation, the T1000 runs Solaris 10 and virtually all Solaris applications. Sun's Web site also lists Gentoo 2006.1 and Ubuntu 6.06 LTS as certified.

The T1000 tested in this article is based on an Ubuntu 6.06 installation. The installation was easy, but required a lot of patience, as the installer obviously is not designed to run on a 9,600bps terminal. Instead of overwriting the current screen with the next, the installation wizard first erases the current screen content, then redraws it

Advertiser Index

For advertising information, please contact our sales department at 1-713-344-1956 ext. 2 or ads@linuxjournal.com. www.linuxjournal.com/advertising

Advertiser	Page #	Advertiser	Page #
Aberdeen, LLC www.aberdeeninc.com	33	LinuxWorld UK www.linuxworldexpo.co.uk	75
AML www.amltd.com	65	Logic Supply, Inc. www.logicsupply.com	93
APPRO HPC Solutions appro.com	C2	LPI www.lpi.org	83
Arcom Control Systems www.arcom.com	15	MBX www.mbx.com	25
ASA Computers www.asacomputers.com	37, 67	Microway, Inc. www.microway.com	C4, 89
Avocent Corporation www.avocent.com/remotecontrol	1	Mikro Tik www.routerboard.com	35
Cari.net www.cari.net	94	Polywell Computers, Inc. www.polywell.com	57
Coraid, Inc. www.coraid.com	23	The Portland Group www.pgroup.com	41
Coyote Point www.coyotepoint.com	3	Qsol.Com www.qsol.com	5
Emac, Inc. www.emacinc.com	87	R1Soft www.r1soft.com	27
EmperorLinux www.emperorlinux.com	21	Rackspace Managed Hosting www.rackspace.com	C3
FairCom www.faircom.com	53	R Cubed Technologies www.rcubedtech.com	55
Genstor Systems, Inc. www.genstor.com	69	SC 2007/Hall-Erickson, Inc. sc07.supercomputing.org	85
HPC Systems, Inc. www.hpcsystems.com	7	Servers Direct www.serversdirect.com	9
Hurricane Electric www.he.net	71	Silicon Mechanics www.siliconmechanics.com	19, 42
Intel www.intel.com	29	Slickedit, Inc. www.slickedit.com	61
Interop New York 2007 www.interop.com/newyork	73	Supermicro www.supermicro.com	39
Iron Systems www.ironsystems.com	79	Swell Software, Inc. www.swellsoftware.com	45
ISPCON www.ispcon.com	70	Technologic Systems www.embeddedx86.com	63
KB Conferences, LLC www.zendcon.com/index.php	91	Uniwide Technologies www.uniwide.com	11
King Star Computer www.kingstarusa.com	31, 77	Verio www.verio.com	47
Linux Journal www.linuxjournal.com	6, 95	White Oak Technologies www.woti.com	13
LinuxCertified, Inc. www.linuxcertified.com	78	Wiley Technology Publishing www.wiley.com	17

Solaris on the T1000

Sun provides several documents with tuning information for Solaris on CoolThreads systems. Linux tuning information, however, is barely available. To check how much impact the lack of tuning options makes, all tests were rerun using Solaris 10 11/06 with the recommended tuning. The bzip2 compression results were virtually the same, although the other benchmarks gained an average of 10%. Whether this 10% stems from the better scalability of Solaris 10 or the extensive tuning is hard to say. However, even with this difference, the T1000 still was far behind the other solutions in most tests.

completely blank and finally, in a third pass, draws the next screen. At 9,600bps, this results in a five-second delay between the screens. Unfortunately, there is no way around this, because in true UNIX spirit, the T1000 does not have a VGA port.

HP Integrity rx2660

The rx2660 is HP's newest low-end Integrity server. It is the first HP Itanium system that shares the chassis with the Proliant line. From the front, it is difficult to distinguish the rx2660 from the 2U DL380G5 without looking at the model number or Intel logo. The rx2660 even has the front VGA port of the DL380—making it the only proprietary system in this review featuring a VGA output.

Like the T1000, the HP server has eight memory slots for up to 32GB of registered DDR2 memory. This is, however, where the similarities end. The rx2660 is a two-socket system and can be equipped with single- or dual-core processors. The single-core processors run at 1.4GHz and offer 6MB of level-three cache. The dual-core processors can be clocked at 1.4GHz (12MB cache) or at 1.6GHz (18MB cache).

Two gigabit Ethernet ports are standard, and the system has eight 2.5" hot-swap SAS drive bays. Depending on which I/O-cage was selected, either three PCI-X slots or one PCI-X and two PCI-Express slots are available for expansion. The server can take a second power supply for redundancy and offers a slot for an optional iLO2 (Integrated Lights-Out 2) remote management card.

Our test system came with two dual-core 1.4GHz CPUs, 4GB of memory and two internal 36GB SAS disks. The iLO2 remote management card was included, bringing the price to $7,095.

The rx2660 is the most versatile unit in this review. It supports HP-UX 11i, OpenVMS v8.3, Windows 2003 and Linux, without changes to the base unit or firmware. HP currently supports Red Hat Enterprise Linux 4 and SUSE Enterprise Server 10. Several other Linux variants, such as Gentoo and Fedora, have Itanium2 versions, but HP currently does not offer support for those flavors.

This rx2660 discussed in this article is based on RHEL 4 Update 4. After powering on the unit, the system starts the EFI firmware. The EFI prompt is menu-based and makes gathering system information and booting the OS very easy. However, after starting the installation from CD, only two lines about the kernel being decompressed are printed. Then, the boot process seemingly stalls. SUSE Enterprise Server showed the same behavior.

An attempt to install HP-UX eventually brought the solution. The system booted normally until "Console is a serial device, no further output will appear on this output device" appeared on the screen. Switching from the VGA port to the serial console worked and allowed RHEL 4 to install without any further issues.

IBM System p5 510Q Express

After changing names several times in the past few years, IBM's Power-based servers are now known under the name IBM System p5. Because of the POWER5 processor's hypervisor, IBM was able to implement the 510Q's most distinguishing feature: LPARs. Short for Logical Partitions, LPARs allow up to 40 OS instances to share the same hardware without the need for any additional software. It even is possible to mix AIX, Red Hat Linux and SUSE Linux on the same server.

The 510Q is equipped with a POWER5+ Quad-Core Module. Due to cooling requirements, the processors in the 510Q are clocked at 1.65GHz—considerably lower than the dual-core model, which comes in 1.9 and 2.1GHz versions. Eight slots can house up to 32GB of DDR2 memory.

Disk storage is provided by up to four internal hot-swap Ultra-320 SCSI drives. Four PCI-X slots are available for expansion. The system also features two gigabit Ethernet controllers.

The back of the system also features two HMC ports. The HMC (short for Hardware Management Console) is a management system that can control up to 254 different LPARs running on up to 48 different servers. Unlike many other p5 models, the 510Q does not require an HMC to operate. Without HMC, the system partitioning capabilities are more limited, but basic features, such as remote console, work without issues.

The p5 510Q used in this review came with four 1.65GHz CPU cores, 6GB of RAM and two 73GB disks. The price was quoted at $6,971.

IBM currently supports AIX 5.2 and 5.3 as well as RHEL 4 and SLES 9 and 10. Gentoo, Fedora and Debian also offer PowerPC distributions. Again, this review is based on the RHEL 4 Update 4. The installation completed without issues and was the easiest installation in this review.

HP Proliant DL140G3

The Proliant DL140G3 is based on Intel's quad-core Xeon 5300 series. This chip essentially is two Core 2 Duo chips mounted on one carrier to fit into a single processor socket. HP has integrated two of these CPUs and up to 16GB of memory into a flat, 1U server. Two disks are available in hot-swap and non-hot-swap variants. The non-hot-swap configuration has space for two expansion PCI-Express slots. In the hot-swap version, one slot is used by an SAS controller. PCI-X variants also are available.

The DL140G3 used in this review was equipped with two Xeon 5345s, 12GB of memory and two hot-swap 36GB SAS disks. The quote came in at $6,531, making the DL140G3 the cheapest server in this comparison.

HP's Web site lists Red Hat Enterprise Linux 3 and 4 as well as SUSE Linux Enterprise Server 9 and 10, all in 32-bit and 64-bit variants. However, none of the 64-bit distributions will boot out of the box. Some searching on the HP Web site led to an advisory recommending disabling the BIOS setting for "8042 Emulation Support". Once the option was turned off, the installation offered no additional surprises.

Growing a World of Linux Professionals

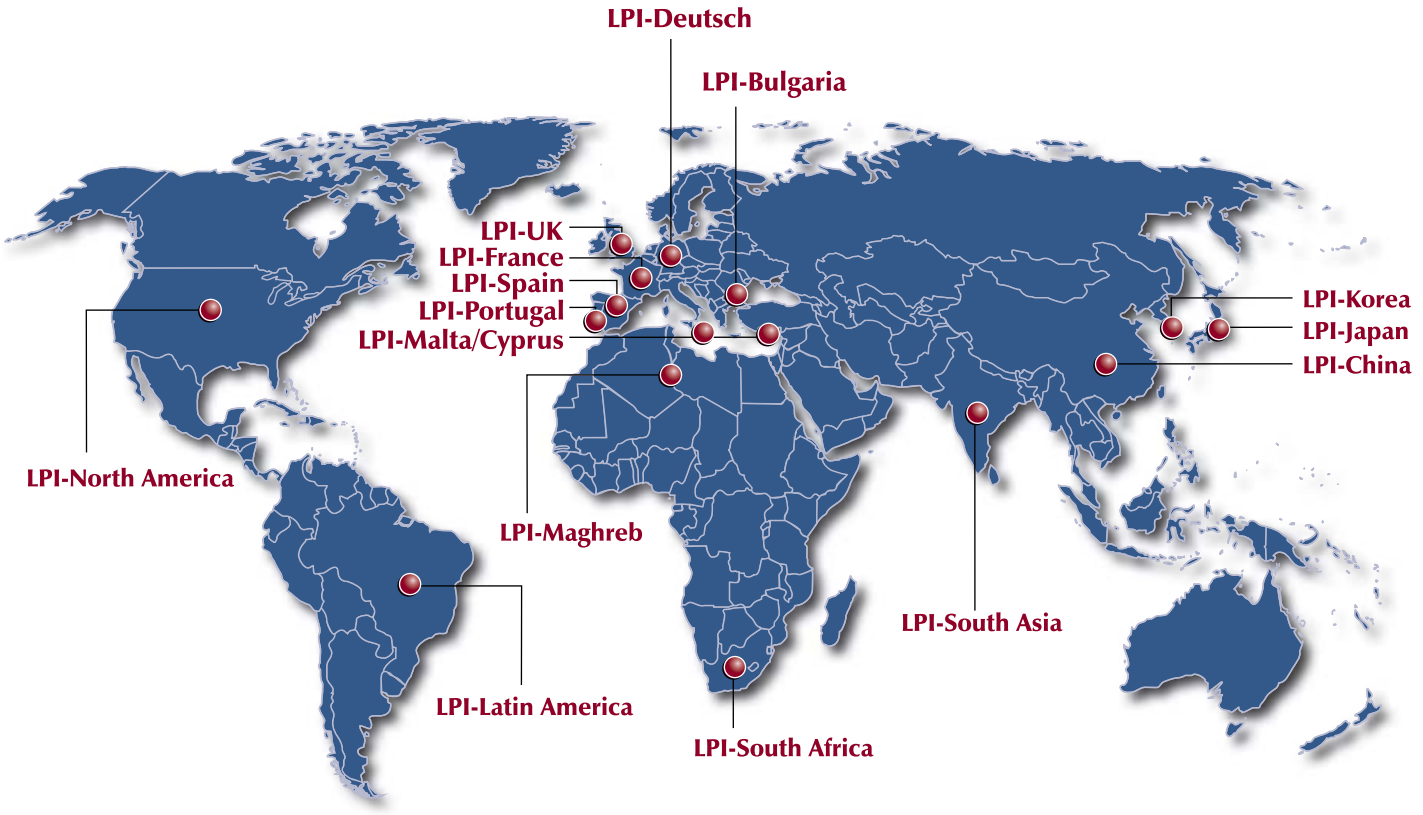

We at the Linux Professional Institute believe the best way to spread the adoption of Linux and Open Source software is to grow a world wide supply of talented, qualified and accredited IT professionals.

We realize the importance of providing a global standard of measurement. To assist in this effort, we are launching a Regional Enablement Initiative to ensure we understand, nurture and support the needs of the enterprise, governments, educational institutions and individual contributors around the globe.

We can only achieve this through a network of local "on the ground" partner organizations. Partners who know the sector and understand the needs of the IT work force. Through this active policy of Regional Enablement we are seeking local partners and assisting them in their efforts to promote Linux and Open Source professionalism.

We encourage you to contact our new regional partners listed above.

Together we are growing a world of Linux Professionals.

Linux Professional Institute

Stable. Innovative. Growing.

The Tests

Reliability and manageability usually are considered the most important features for the proprietary systems. However, in recent years, management capabilities have increased on the x86-based servers. At the same time, the low-end systems in this comparison have lost many of these features their big brothers have. As an example, Sun's T1000 does not even provide hot-swappable disks.

For this reason, the tests in this article focus on performance, and the systems have to prove themselves in five different scenarios.

File Compression

File compression is a CPU-intensive task with very low I/O requirements. The first test was run with a single bzip2 -1 (lowest compression) process compressing a 2GB file. This established the baseline performance for each system. Then the test is rerun with 2, 4, 8, 16 and 32 concurrent processes compressing the same 2GB file as before. These additional processes allow the systems to use more of the available processor resources. Because the processes are independent, scaling should be as close to linear as the hardware allows.

After the first run, all benchmarks were executed a second time at the highest compression level, -9. As the man page describes, the higher compression level significantly increases the memory usage of the process.

The most interesting result in this test is the T1000. Just as Sun stated, the single-thread performance of the CPU is very weak. However, once 32 threads are being executed simultaneously, the system beats the rx2660.

The second interesting result is the DL140. As soon as eight bzip2 -9 threads are executed, the cache (4MB shared between each two cores) is no longer able to contain all data required. The performance hit is substantial. Although at low concurrency, the difference between low and high compression is below 10%, at 32 threads, the difference is 111%. The other systems show almost the same performance with both compression levels.

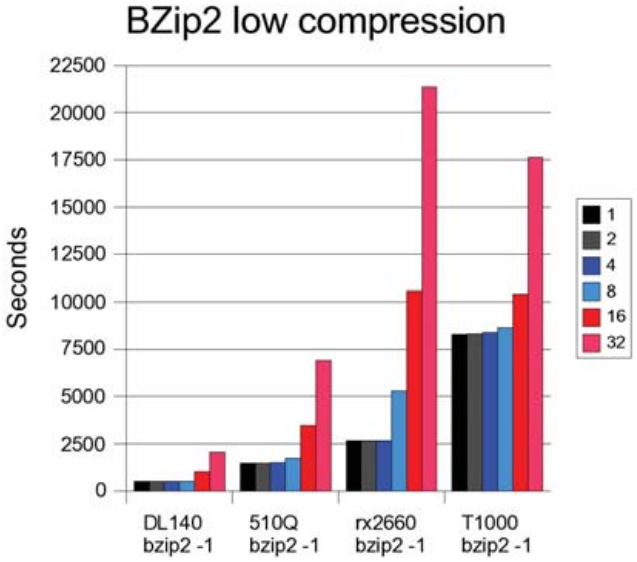

Figure 1. Comparison of bzip2 Low Compression Performance

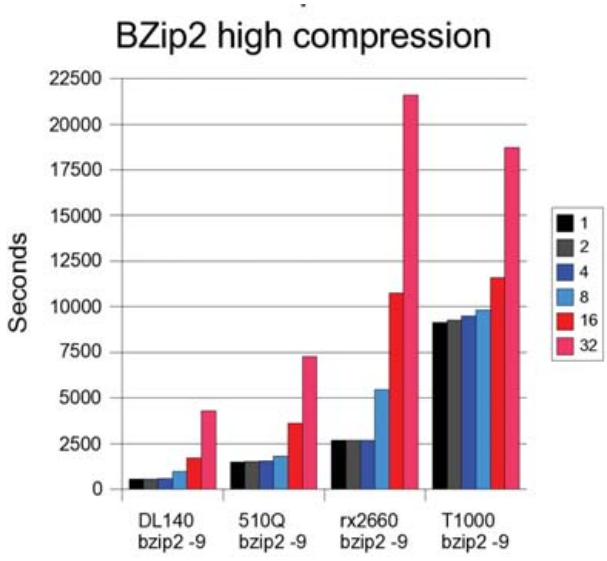

Figure 2. Comparison of bzip2 High Compression Performance

C++ Compilation

As with file compression, compiling C++ code is another scenario with high CPU use and low demands on the I/O and memory subsystems. The major difference, however, is that the compiler instances are not independent. The way most C++ projects lay out their makefiles allows the make program to kick off compiles in only one directory at a time. This limits the number of compiler processes that can be started.

Also, several portions of the build, like dependency generation and linking, cannot be parallelized at all. This makes the C++ compiler test much less thread-friendly but more realistic.

The subject of this test was the Perl 5.8.8 source code. Configure was run accepting all defaults except the library path (/usr/lib64 was missing on the Xeon system), and the optimization setting was increased from the default -O2 to -O4. The compiles were run with

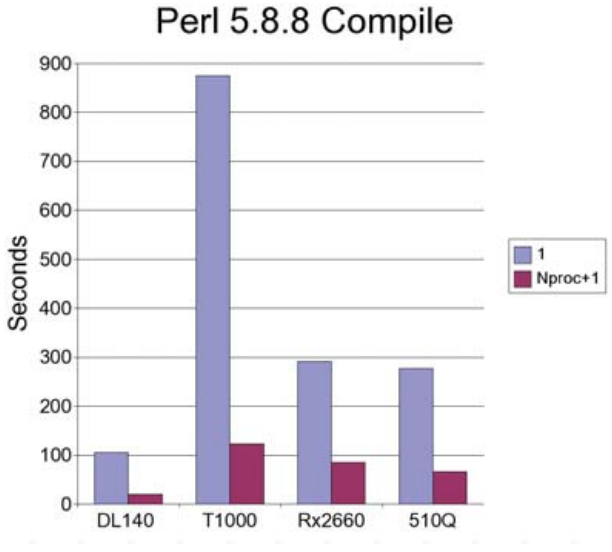

Figure 3. Comparison of Performance with Perl 5.8.8 Compilation

SC07

Reno-Sparks
Convention Center
Reno, Nevada

November 10-16, 2007

SC07, the premier international conference on high performance computing, networking, storage, and analysis, will convene in November 2007 in Reno, Nevada. SC07 expands the frontiers of those high performance technologies with a new initiative in high performance clusters. Of course, SC07 will also host the exceptional technical program, tutorials, workshops, exhibits, education program, and other exciting activities you've come to expect with the SC conferences.

SC is the one place attendees can see tomorrow's technology being used to solve world-class challenge problems today. This series has become a unique gathering spot for international scientists, researchers, software developers, network engineers, policy makers, corporate managers, CIOs, and IT administrators. SC07 represents the latest and greatest in technology, applications, industrial products, research results, and national/international initiatives. *We look forward to your participation!*

Conference Location:
Reno, Nevada

Conference Dates:
November 10-16, 2007

Exhibition Dates:
November 13-15, 2007

For more information, visit:
http://sc07.supercomputing.org

Sponsors:
ACM SIGARCH/IEEE Computer Society

acm | IEEE | IEEE COMPUTER SOCIETY

To compete with the established x86 solutions and the massive budget Intel can invest into CPU development, IBM, Sun and the Intel Itanium team have to be innovative and take ideas to new heights.

one thread and then with one thread more than the number of available CPUs.

The results were similar to the compression benchmark. Again, the T1000 profited the most from the additional threads, but even at the highest settings, it was not able to keep up with the other solutions.

MySQL Database

MySQL is, without question, the best known open-source database; however, its scalability has been questioned on many occasions. Especially in environments that have a larger percentage of writes to the database, the performance is said to suffer in larger SMP systems. This means that systems that rely on a large number of threads have a disadvantage, and systems with high single-core performance should fare better.

The exact version of MySQL depends on the distribution used. Red Hat Enterprise Linux 4 includes MySQL 4.1.20. The T1000 running Ubuntu 2006.6 LTS was running the much newer version 5.0.20. Comparing such different versions sounds strange, but it is in the spirit of the article—compare the servers the way they come and are supported by the vendors. In most enterprise environments, compiling your own version of MySQL is simply not an option—something that is especially painful for the Itanium-based system. To provide a better comparison, the T1000 also was tested with MySQL 4.1.20.

To test MySQL performance, Sysbench 0.4.8 was used. Sysbench is designed to create a workload that is similar to an OLTP load in a real system. The exact command run was:

```
sysbench --test=oltp --num-threads=512 --mysql-user=root
➥--max-time=240 --max-requests=0
```

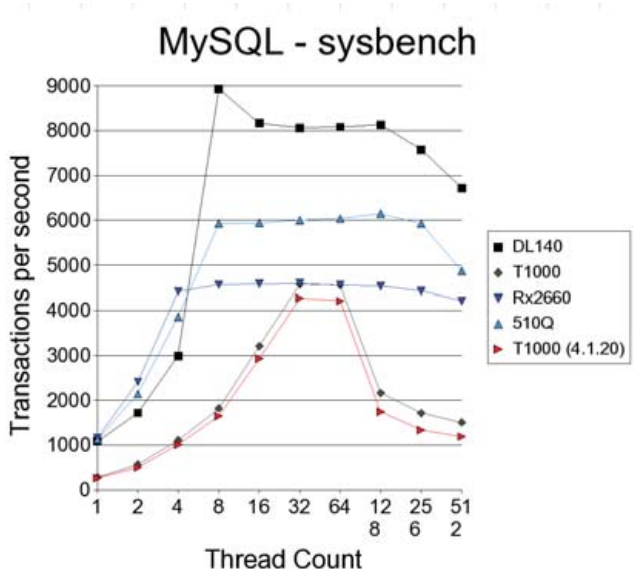

Figure 4. The DL140 benefits most from multiple threads.

The most interesting result in this test was the rx2660. Although all other systems showed a larger performance decrease when being tested with a large thread count, the Itanium system managed to keep virtually the same performance numbers under load.

PostgreSQL Database

PostgreSQL is another open-source database. It is not as widespread as MySQL, but many comparisons show that PostgreSQL has better scalability, because of the row version mechanism (MVCC) used. Red Hat shipped PostgreSQL 7.4.16, and Ubuntu came with 8.1

Because Sysbench requires PostgreSQL 8.0 or newer, the tool used to benchmark PostgreSQL was pgbench. The scaling factor selected was 50. Because pgbench results vary greatly, the tests were rerun 32 times for each number of clients and the highest result was taken.

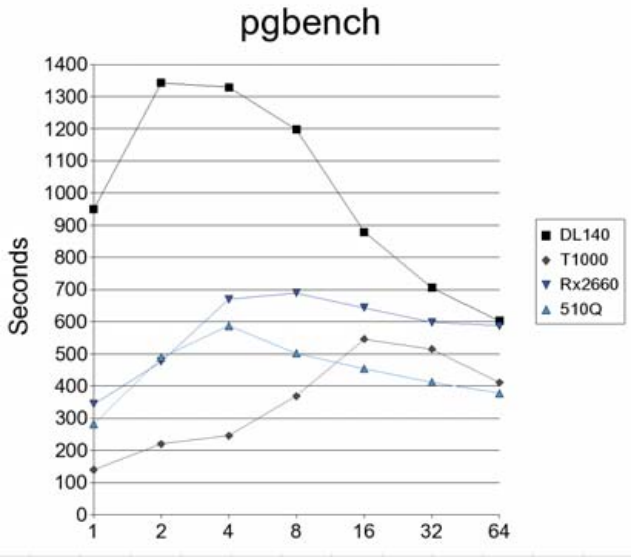

Figure 5. Comparison of Performance Using the PostgreSQL pgbench

The PostgreSQL benchmarks look much like the MySQL results before. Notice, however, the large drop-off of the Xeon system compared with the other systems. The T1000, however, profited from the better scalability of PostgreSQL.

Web Server—PHP Application

The execution of PHP scripts combines CPU, memory and disk usage. For testing purposes, a small PHP script was written that executes a few MySQL database queries and formats the output into very simple HTML. Additional CPU load stems from compilation of the script (no PHP accelerator was used) and a loop in the middle of the script. An fopen call to a random file and a fread of the first kilobyte was used to simulate disk access.

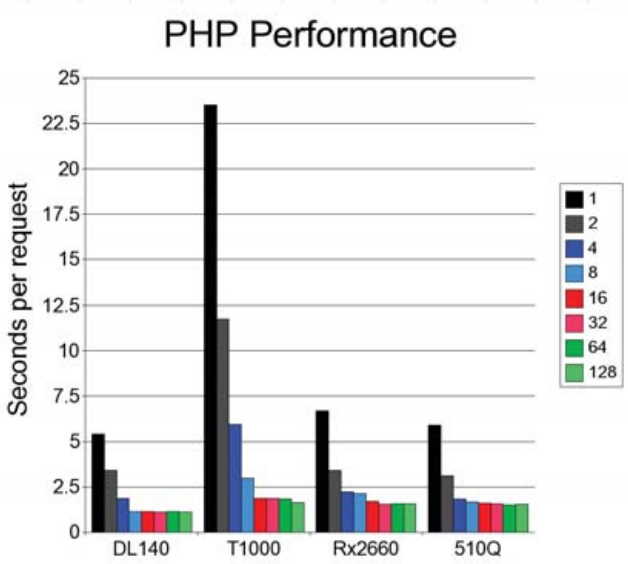

Figure 6. PHP application performance shows mostly narrow differences.

In this benchmark, the performance gap between the different solutions was much more narrow than before. When fully utilized, the three proprietary solutions performed similarly. The T1000 was only a few percentages slower than the POWER5 and Itanium systems. The Xeon, however, maintained at least a 35% lead throughout the test.

Compatibility

Because the tests in this article are all based on open-source software, no compatibility issues were observed. Of course, low-level software that accesses hardware directly has to be customized for the different systems, but all distributions were feature-complete and included all common programs for both desktop and server use.

Once you look at closed-source software, the picture unfortunately changes. The Itanium processor is fairly well supported, while most software that supports the Power platform comes directly from IBM. Worse off is the T1000. Not even the Java JDK is available from Sun.

Conclusion

Although the T1000 consistently came in last, it looked better as the more threads were working concurrently. However, because most Linux developers are using single processor or dual-core systems, it is hard to find open-source applications that are capable of starting 32 threads at once.

The third place goes to the Itanium-based rx2660. The Itanium processor performed well on single-threaded applications, but in the end, it was beat consistently by the POWER5-based 510Q. With an improved version of GCC, Intel and HP surely could change this picture, but for now, there is little chance that the distributions will adopt a proprietary compiler to gain performance.

Eight execution threads earned the IBM System p5 510Q the second place in this comparison. The 510Q bested the T1000 and also held a consistent lead over the rx2660 once all eight threads were utilized. In addition, the possibilities of partitioning the system without the use of Xen or VMware makes this system the best choice among the proprietary boxes.

The biggest surprise, however, was the DL140G3. Originally, it was planned only as a point of reference, but Intel has designed a very impressive solution with the latest quad-core Xeons. For years, Intel or AMD systems running Windows or Linux have competed well against smaller UNIX systems, but never before has an x86-based system enjoyed a performance lead like this. In addition, HP has done an excellent job integrating management capabilities into the server.

In one sentence—there is little to no reason to go with the low-end proprietary server. Performance is worse, and at the low end, reliability features are comparable. Does that mean these chips are dead? Not by a long shot. Intel or AMD systems usually don't go beyond 16 cores, while the UNIX vendors offer systems with up to 144 cores. However, most of these large systems offer no or limited Linux support. In addition to the high CPU count, the virtualization capabilities of the POWER5 systems are impressive—low overhead at no additional cost.■

Peter Arremann currently works at Verizon on system administration, automation of software development processes and the company's open-source policy. He has been using Linux for more than a decade and can be reached at loony@loonybin.org.

Introduction to Stream Control Transmission Protocol

The first in a series on SCTP and how to migrate TCP applications.

JAN NEWMARCH

Most people who have written networking software are familiar with the TCP and UDP protocols. These are used to connect distributed applications and allow messages to flow between them. These protocols have been used successfully to build Internet applications as we know them: e-mail, HTTP, name services and so forth. But, these protocols are more than 20 years old, and over time, some of their deficiencies have become well known. Although there have been many attempts to devise new general-purpose transport protocols above the IP layer, only one so far has received the blessing of the IETF: SCTP (Stream Control Transmission Protocol). The central motivation behind SCTP is to provide a more reliable and robust protocol than either TCP or UDP that can take advantage of features such as multihoming.

SCTP is not a radical departure from TCP or UDP. It borrows from both but is most similar to TCP. It is a reliable session-oriented protocol, like TCP. It adds new features and options and allows finer control over the transport of packets. In all but the "edge" cases, it can be used as a drop-in in place of TCP. This means that TCP applications often can be ported trivially to SCTP. Of course, to benefit properly from the new features of SCTP, you need to use the additional API calls for SCTP.

The first additional feature in SCTP is better support for multihomed devices—that is, computers with more than one network interface. At one time this meant only routers and bridges connecting different parts of the Internet, but now even computers on the edges of the network can be multihomed. Most laptops have built-in Ethernet cards and Wi-Fi cards, and many have Bluetooth cards as well (which have IP support through the Bluetooth PPP stack). Some laptops now are shipping with WiMAX cards, and it even is possible to run IP over the infrared port! So, the standard laptop is at least dual-homed, with possibly up to five distinct IP network interfaces.

TCP and UDP allow use of only one or all of the interfaces. But, what if you are running your laptop as a peer in, say, a file-sharing service? It probably would be silly to use the Bluetooth and infrared interfaces. WiMAX can be very expensive to shift large amounts of data. But, it would make sense to use both the Ethernet and Wi-Fi interfaces. SCTP can support this selective choosing of interfaces. Some implementations even can add and drop interfaces dynamically, so as you unplug your laptop and move out of the house, an application can switch to the WiMAX interface if you want.

The second main new feature is multistreaming—that is, one "association" (which is renamed from "connection" from TCP) can support multiple data streams. It is no longer necessary to open up multiple sockets; instead, a single socket can be used for multiple streams to a connected host. Several TCP applications could benefit from this. For example, FTP (the major file transfer protocol) uses two streams: one on port 21 for control messages and another on port 20 for data. This caused problems with firewalls in place. A client could connect to a server through a firewall, but the server could not connect to the client for data transfer because of the firewall. The FTP protocol had to be extended to allow for "passive" connections to overcome this. There would be no need for such an extension under SCTP—simply send the data on a separate stream in an association established by a client.

The X Window System also uses multiple sockets on multiple ports. Although it is not common, a computer can have multiple display devices. Typically, the first is on port 6000, the second on port 6001 and so on. Under SCTP, these could all be separate streams on a single association. HTML documents often contain embedded references to image files, and to display a page properly requires downloading the original page and all of these images (or embedded frames too). HTTP originally used a separate TCP connection per downloaded URL, which was expensive and time consuming. HTTP 1.1 brought in "persistent connections", so that a single socket could be reused for all of these sequential downloads. Under SCTP, the separate images could be downloaded concurrently in separate streams on a single association.

There are even more subtle uses of SCTP multiple streams. An MPEG movie consists of different types of frames: I frames, P frames and B frames. I frames encode complete images, and the other two types measure differences between frames. Typically, there is an I frame every ten frames, with the others "predicted" from these. It is critical that the I frames be delivered, but less so for the P and B frames. Although SCTP is not designed as a Quality-of-Service protocol, it does allow different delivery parameters on different streams within an association, so that the I frames can be delivered more reliably.

SCTP has many more features, such as:

- TCP is a byte-oriented protocol, and UDP is message-oriented. The majority of applications are message-oriented, and applications using TCP have to jump through hoops, such as sending the message length as a first parameter. SCTP is message-oriented, so such tricks are not so necessary.

- A single socket can support multiple associations—that is, a computer can use a single socket to talk to more than one computer. This is not multicast, but it could be useful in peer-to-peer situations.

- SCTP has no "out of band" messages, but a large number of events can be interleaved onto a single association, so that an application

Affordable InfiniBand Solutions

4 Great Reasons to Call Microway NOW!

1 TriCom™

- DDR/SDR InfiniBand HCA
- "Switchless" serial console
- NodeWatch web enabled remote monitor and control

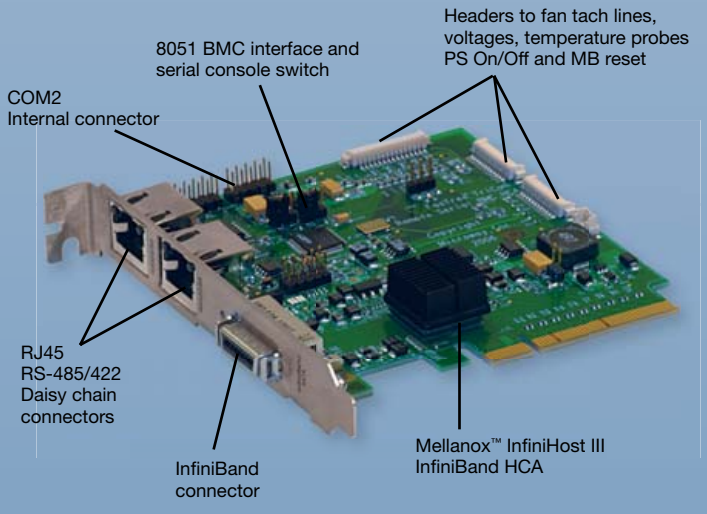

2 FasTree™

- DDR InfiniBand switches
- Low latency, modular design
- 24, 36 and 48 port building blocks

3 ServaStor™

- Extensible IB based storage building blocks
- Redundant and scalable
- Parallel file systems
- Open source software
- On-line capacity expansion
- RAID 0,1,1E, 3, 5, 6, 10, 50

4 InfiniScope™

- Monitors ports on HCA's and switches
- Provides real time BW diagnostics
- Finds switch and cable faults
- Lane 15 interface
- Logs all IB errors

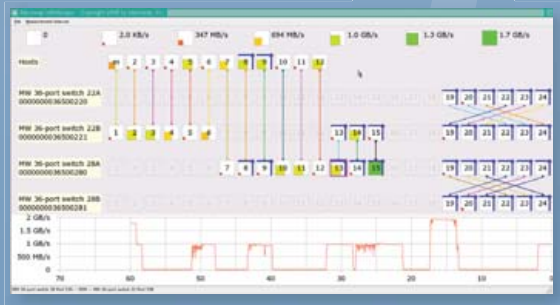

Upgrade your current cluster, or let us design your next one using Microway InfiniBand Solutions.

To speak to an HPC expert call ***508 746-7341*** *and ask for technical sales or email*
sales@microway.com
www.microway.com

Microway®
Technology you can count on℠

Listing 1. echo_client.c

```
#define USE_SCTP

#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>

#ifdef USE_SCTP
#include <netinet/sctp.h>
#endif

#define SIZE 1024
char buf[SIZE];
char *msg = "hello\n";
#define ECHO_PORT 2013

int main(int argc, char *argv[]) {
        int sockfd;
        int nread;
        struct sockaddr_in serv_addr;
        if (argc != 2) {
                fprintf(stderr, "usage: %s IPaddr\n", argv[0]);
                exit(1);
        }
        /* create endpoint using TCP or SCTP */
        sockfd = socket(AF_INET, SOCK_STREAM,
#ifdef USE_SCTP
                        IPPROTO_SCTP
#else
                        IPPROTO_TCP
#endif
                );
        if (sockfd < 0) {
                perror("socket creation failed");
                exit(2); }
        /* connect to server */
        serv_addr.sin_family = AF_INET;
        serv_addr.sin_addr.s_addr = inet_addr(argv[1]);
        serv_addr.sin_port = htons(ECHO_PORT);
        if (connect(sockfd,
                    (struct sockaddr *) &serv_addr,
                    sizeof(serv_addr)) < 0) {
                perror("connect to server failed");
                exit(3);
        }
        /* write msg to server */
        write(sockfd, msg, strlen(msg) + 1);
        /* read the reply back */
        nread = read(sockfd, buf, SIZE);
        /* write reply to stdout */
        write(1, buf, nread);

        /* exit gracefully */
        close(sockfd);
        exit(0);
}
```

can monitor the state of the association (for example, when the other end adds another interface to the association).

- The range of socket options is greater than TCP or UDP. These also can be used to control individual associations or individual streams within a single association. For example, messages on one stream can be given a longer time-to-live than messages on other streams, increasing the likelihood of their delivery.

Availability of SCTP

The SCTP Web site (**www.sctp.org**) has a list of implementations of SCTP. There are implementations for BSD and Windows, and since 2001, there has been a Linux kernel project at **sourceforge.net/projects/lksctp**. At present, SCTP is not in any Microsoft release, so applications running on Windows need to install one of the available stacks.

SCTP is included in the Linux kernel as an experimental network protocol. SCTP is normally built as a module. It may be necessary to load the module using `modprobe sctp`. To build user applications, you may need to install the SCTP tools—in Fedora Core 6, these are in the RPM packages lksctp-tools-1.0.6-1.fc6.i386.rpm and lksctp-tools-devel-1.0.6-1.fc6.i386.rpm. On Fedora Core 6, I also had to add a symbolic link from /usr/lib/libsctp.so to /usr/lib/libsctp.so.1.

The lksctp-tools package contains the libraries to run SCTP applications. It also contains a program called checksctp, which tells you if your kernel has support for SCTP. When you run this program, it prints either "SCTP supported" or an error message.

The devel package contains the sctp.h header file, so you can compile and build your own applications, and man pages for the SCTP function calls.

Firewalls

Most firewalls can be configured to deal with SCTP packets, but the documentation for each firewall may not mention SCTP explicitly. For example, the man page for iptables says, "The specified protocol [in a rule] can be one of tcp, udp, icmp, or all...". But, it then goes on to say, "A protocol name from /etc/protocols is also allowed", and in that file, we find that protocol 132 is sctp. So, rules for SCTP can be added to iptables in the same way as TCP and UDP rules.

For example, an iptables rule to accept SCTP connections to port 13 would be:

```
-A INPUT -p sctp -m sctp -i eth0 --dport 13 -j ACCEPT
```

Webmin is a popular administration tool for managing things like iptables rules. Unfortunately, as of version 1.340, it could not

accept this rule, because it is hard-wired to accept port numbers only for TCP and UDP, not realising that SCTP also uses port numbers. Such a rule would need to be entered by hand into the iptables configuration file /etc/sysconfig/iptables. This will be fixed in later versions of Webmin after I logged a bug report, but similar problems may occur in other tools.

One-to-One Socket API

As with TCP and UDP, SCTP provides a socket API for applications. A server creates a socket bound to a port and then uses this to accept a connection from a client. A client also creates a socket and then connects to a server. Both then use the socket file descriptor to read and write messages. SCTP is not a superset of TCP. Nevertheless, when restricted to a similar style of connection as TCP, there are sufficient similarities that an SCTP socket often can be used as a drop-in replacement for a TCP socket. When used in this way, SCTP sockets are called one-to-one sockets, as they simply connect one host to a single other host.

To create a TCP socket, use the system call:

```
sockfd = socket(AF_INET, SOCK_STREAM, IPPROTO_TCP)
```

This creates an IPv4 socket. To create an IPv6 socket, replace the first parameter with AF_INET6. The last parameter often is given as zero, meaning "use the only protocol value in the family". It is better to use IPPROTO_TCP explicitly, because SCTP introduces another possible value.

To create an SCTP one-to-one socket, simply replace IPPROTO_TCP with IPPROTO_SCTP:

```
sockfd = socket(AF_INET, SOCK_STREAM, IPPROTO_SCTP)
```

and that (in many cases) is it! The client or server is now talking the SCTP protocol instead of TCP.

To see this in action, Listings 1 (echo_client.c) and 2 (echo_server.c) give a simple echo-client and server, where the server returns a string sent to it when a client connects to it. Only the line above needs to change in both the client and the server (with also an extra include file, sctp.h).

The usual C compile command can be used to create object modules and executables. If the program uses SCTP-specific functions (the programs in Listings 1 and 2 don't), you also need to link in the SCTP library:

```
cc -o echo_client echo_client.c -lsctp
```

Is it worthwhile to take an application that runs over TCP and move it to SCTP? The disadvantages are that SCTP is not as well supported as TCP, the tools are sometimes not aware of SCTP and the API is still evolving. On the other hand, it benefits from the experience of 20 years of seeing TCP and UDP applications in practice. For example, SCTP is secure from SYN attacks by design, and the protocol has no known security holes. SCTP also will take advantage of multihoming when needed automatically. If packets are getting lost, due to, say, congestion, SCTP will use different interfaces to try to avoid the losses, and this could result in faster throughput.

The central motivation behind SCTP is to provide a more reliable and robust protocol than either TCP or UDP that can take advantage of features such as multihoming.

The withsctp Tool

In the previous section, I discussed how to alter the source code of a client or server to use SCTP instead of TCP. The sctp-tools package contains a program called withsctp, which essentially does the same to binary code. This program acts as a wrapper around a TCP application to turn it into an SCTP application. It first saves the address of the "real" socket() function call, and then inserts its own version of socket() into the load library path. This new version of socket() simply gets the parameters of the function call, changes the third parameter from IPPROTO_TCP to IPPROTO_SCTP and calls the "real" socket() function.

For example, the xinetd dæmon can run a group of TCP and UDP

Focused on the fastest growing web development language, the **2007 Zend/PHP conference** will bring together PHP practitioners and business managers from around the world for three days of exceptional presentations and outstanding networking events.

Join PHP experts and community members who will share their experiences and best practices in more than 40 sessions from beginner level to leading-edge advanced topics.

Zend
The php Company

Please visit our website and register online.
www.zendcon.com

Listing 2. echo_server.c

```
#define USE_SCTP

#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>

#ifdef USE_SCTP
#include <netinet/sctp.h>
#endif

#define SIZE 1024
char buf[SIZE];
#define ECHO_PORT 2013

int main(int argc, char *argv[]) {
        int sockfd, client_sockfd;
        int nread, len;
        struct sockaddr_in serv_addr, client_addr;

        /* create endpoint using TCP or SCTP */
        sockfd = socket(AF_INET, SOCK_STREAM,
#ifdef USE_SCTP
                        IPPROTO_SCTP
#else
                        IPPROTO_TCP
#endif
                );
        if (sockfd < 0) {
                perror("socket creation failed");
                exit(2);
        }
        /* bind address */
        serv_addr.sin_family = AF_INET;
        serv_addr.sin_addr.s_addr = htonl(INADDR_ANY);
        serv_addr.sin_port = htons(ECHO_PORT);
        if (bind(sockfd,
                  (struct sockaddr *) &serv_addr,
                  sizeof(serv_addr)) < 0) {
                perror("bind failed");
                exit(3); }
        /* specify queue length */
        listen(sockfd, 5);
        for (;;) {
                len = sizeof(client_addr);
                /* get a connection from client */
                client_sockfd = accept(sockfd,
                                (struct sockaddr *) &client_addr,
                                &len);
                if (client_sockfd == -1) {
                        perror("accept failed");
                        continue;
                }
                /* transfer data */
                nread = read(client_sockfd, buf, SIZE);
                /* write to stdout */
                write(1, buf, nread);
                /* and echo it back to client */
                write(client_sockfd, buf, nread);
                /* no more for this client */
                close(client_sockfd);
        }
}
```

services. The services are those listed in the directory /etc/xinetd.d, which have `enable = yes` or `disable = no`. The TCP services all can be run over SCTP by:

```
withsctp xinetd
```

One of the simplest services that is run by xinetd is daytime. The service accepts a connection and returns an ASCII string for the current date. A quick Google search turns up source code for many clients, but the simplest way is to run Telnet:

```
telnet <host-name> 13
```

If you have daytime running as an SCTP service rather than a TCP service, use withsctp to connect to it:

```
withsctp telnet <host-name> 13
```

This is a quick way of testing whether a TCP service can be converted to SCTP.

Message Orientation

TCP is a byte-oriented protocol—that is, you write bytes and read bytes. The UNIX system calls read() and write() typically are used for this. TCP also has send()/recv(), which have an extra flags parameter, but these do not change the byte-transfer model.

SCTP, on the other hand, is message-oriented, more like UDP. Most Internet applications have a message structure to their communications rather than merely a sequence of bytes. For example, a single HTTP request has a header and body section, and even the header section is composed of an arbitrary number of lines. The sender has to compose the parts into the single request, and the receiver of such a message has to parse it back into its component messages. A few protocols are only byte-oriented (for example, the file transfer mode of FTP), but these are the minority.

SCTP makes it easy to use a message-based structure—within limits. A write() call writes a complete message. The corresponding read() reads this complete message. So, to send an HTTP header over SCTP, you could do a write of each line, followed by a write of an empty line. The receiver would read each line as a separate message, stopping after reading an empty line. There would be no need

Listing 3. read_sctp_msg.c

```
#include <stdio.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <stdlib.h>
#include <netinet/sctp.h>

/* call by
     nread = read_sctp_msg(sockfd, &msg)
*/
int read_sctp_msg(int sockfd, uint8_t **p_msg) {
        int rcv_buf_size;
        int rcv_buf_size_len = sizeof(rcv_buf_size);
        uint8_t *buf;
        struct sockaddr_in peeraddr;
        int peer_len = sizeof(peeraddr);
        struct sctp_sndrcvinfo sri;
        int total_read = 0;

        *p_msg = NULL; /* default fail value */

        if (getsockopt(sockfd, SOL_SOCKET, SO_RCVBUF,
                        &rcv_buf_size, &rcv_buf_size_len) == -1) {
                return -1;
        }
        if ((buf = malloc(rcv_buf_size)) == NULL) {
                return -1;
        }

        while (1) {
                int nread;
                int flags;

                nread = sctp_recvmsg(sockfd, buf+total_read,rcv_buf_size,
                                        (struct sockaddr *) &peeraddr,&peer_len,
                                        &sri, &flags);
                if (nread < 0) {
                        return nread;
                }
                total_read += nread;

                if (flags & MSG_EOR) {
                        /* trim the buf and return msg */
                        printf("Trimming buf to %d\n", total_read);
                        *p_msg = realloc(buf, total_read);
                        return total_read;
                }
                buf = realloc(buf, total_read + rcv_buf_size);
        }

        /* error to get here? */
        free(buf);
        return -1;
}
```

to parse the received bytes into a set of lines before processing each one. Note that if the original TCP application already used a series of writes followed by a single read, expecting TCP to concatenate all the messages, the application would need to be modified to match each write to a corresponding read statement.

The caveats are with big messages. Applications that want to take advantage of these messaging capabilities must be careful when sending big messages (say 32KB or more). To send a message, you aren't merely passing a pointer to data on the stack, you're actually moving that data across the network. That means putting it into buffers on the sender side, passing it through buffers in intermediate nodes and, finally, delivering it to a buffer in the reading application. All of these buffers have limits that cannot be exceeded.

For example, say a sender uses a buffer with a size set by the socket option SO_SNDBUF. An attempt to write a message larger than that will fail and return -1. The size of this is generous, typically about 64KB. It can be changed by using `setsockopt(sockfd, SOL_SOCKET, SO_SNDBUF, &val, &val_len)`, where val is an integer variable containing the length to which you want to set the buffer. But, then other limits may come into play. Each host along the route from sender to receiver will have a maximum packet size that it will pass along. The Path Maximum

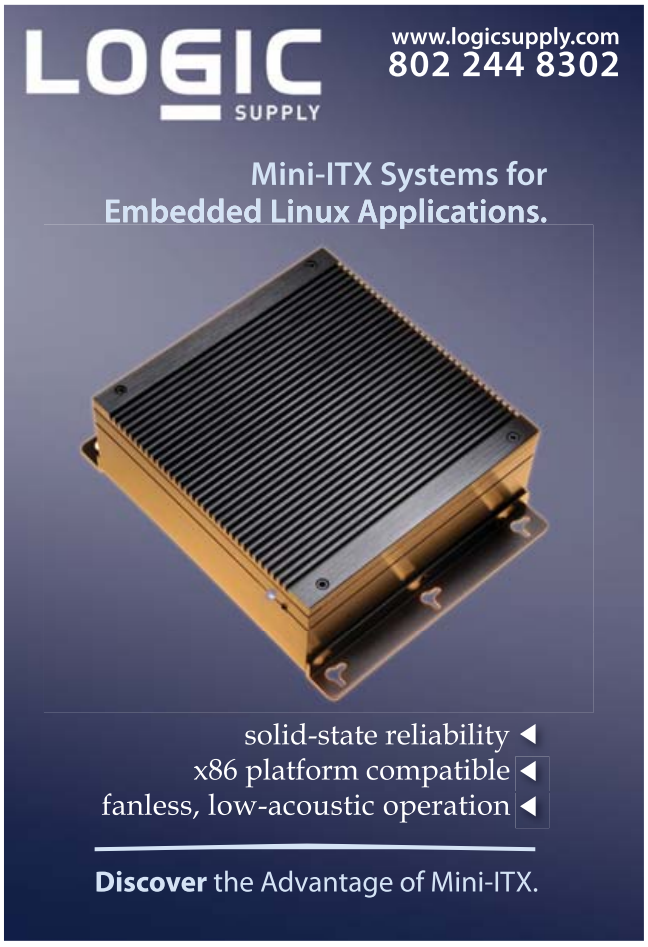

It is no longer necessary to open up multiple sockets; instead, a single socket can be used for multiple streams to a connected host.

Transmission Unit (PMTU) is the minimum of all of these. If the message (plus any IP and SCTP headers) is larger than the PMTU, it will be fragmented and delivered in pieces. The sender can guard against this by setting the SCTP option SCTP_DISABLE_FRAGMENTS so that a message is delivered as a single entity or not at all, but this typically will only decrease the maximum possible message size.

The receiver of a message also has a receiving buffer size, which is controlled by the socket option SO_RCVBUF. It will not receive messages larger than this—fragmenting them if necessary. The major problem from the receiving side is how to deal with fragmented messages. The system calls read() and recv() do not contain any information about message boundaries, as they are byte-oriented. Fortunately, SCTP has a new system call, sctp_recvmsg(), which returns status information about the read in an integer parameter. In particular, if the MSG_EOR bit (message end-of-record) is set, read of a message has been completed. If it is not set, the message has been fragmented and more of the message needs to be read. This can be used by the reader to build up a complete message before processing it.

Listing 3 shows how the sctp_recvmsg() call can be used to receive fragmented messages and build them up into a complete message. It does so by reading each part of a message as it comes in and adding it to the parts already received. When a part arrives with the MSG_EOR bit set in the flags, the message is complete and can be returned to the reading application.

IPv6

SCTP has full support out of the box for IPv6 as well as IPv4. You simply need to use IPv6 socket addresses instead of IPv4 socket addresses. If you create an IPv4 socket, SCTP will deal only with IPv4 addresses. But, if you create an IPv6 socket, SCTP will handle both IPv4 and IPv6 addresses.

Conclusion

This article provides a brief introduction to the IETF Stream Control Transmission Protocol and explains how it can be used as a replacement for TCP. In future articles, we will examine additional features of SCTP and show their use.■

Jan Newmarch is Honorary Senior Research Fellow at Monash University. He has been using Linux since kernel 0.98. He has written four books and many papers and also has given courses on many technical topics, concentrating on network programming for the last six years. His Web site is jan.newmarch.name.

Resources

The Principal Site for SCTP (contains pointers to the RFCs and Internet Drafts for SCTP): **www.sctp.org**

The Linux Kernel Project Home Page: **https://lists.sourceforge.net/lists/listinfo/lksctp-developers**.

Stream Control Transmission Protocol (SCTP): A Reference Guide by Randall Stewart and Qiaobing Xie, Addison-Wesley.

Unix Network Programming (volume 1, 3rd ed.) by W. Richard Stevens, et al., has several chapters on SCTP, although some of it is out of date.

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

The Ultimate Linux PVR

After much wasted time trying to re-create TiVo, it turns out the real thing beats all.

Nick Petreley, Editor in Chief

You may recall that, quite some time ago now, I started a MythTV project. The idea was to create a Personal Video Recorder (PVR) based on Linux that could handle HDTV. I mistakenly assumed that it would not only be possible to create the ideal HDTV PVR on Linux, but also that it would be a relatively simple task for a hacker like myself. The project consumed a great deal of my time for a few months, most of which turned out to be wasted.

The bottom line is that cable companies make it nearly impossible to get decent quality HDTV except through their own cable boxes. HDTV is under the control of Digital Rights Management (DRM) robot overlords, and the overlords have spoken. They will not risk the possibility that we might record something like an HDTV movie on our personal computers. Theoretically, at least, it's not entirely impossible to route around the roadblocks. You may be able to get a cable box with an output you can capture (like IEEE 1394, for example). But, you won't be able to decode HDTV cable signals from the coax, and as far as I can tell, it isn't yet possible to connect the HDMI output of a cable box to a PC in a way that is useful to Linux.

If you're satisfied with standard definition TV, MythTV will do the job for you. I had few problems getting MythTV and a variety of capture cards working. I tried the Hauppauge WinTV PVR-500, the pcHDTV HD-5500 and the DVICO FusionHDTV5 Gold. They all worked more or less, but they were all overkill. The HDTV cards were overkill, because I could use them only in standard definition. The PVR-500 was overkill, because the best configuration used the SDTV output of the cable box, and there was no need for the two tuners on the card.

MythTV is not all that hard to set up, but it makes you work harder than you should. For example, you need to change some scan settings to get rid of blinking on-screen caused by the closed-caption signal. This sort of thing should be automatic, and it probably will be eventually (if the MythTV folks haven't already done this since I tried it).

If you don't want to get your hands so dirty with MythTV, the commercial product SageTV (**www.sagetv.com**) is a very nice alternative. Whichever you use, stick with a cheap Hauppauge card like the WinTV PVR-150. It is really all you need if you are going to combine your PVR with a cable box. The only reason to get a dual-tuner card is if you want to use it with unscrambled cable or want to use over-the-air TV.

The first alternative I tried was an HDTV-capable cable box with built-in PVR. What a disaster. They're not intelligent at all about getting the equivalent of a TiVo season pass, where the PVR records shows by name instead of specific time slots. The fast-forward and rewind features lag seconds behind when you press the button, and it's almost impossible to navigate through a recorded show. I always would overshoot the destination and have to watch more of the show than I wanted.

I finally invested in the Linux-based TiVo Series 3 with two cable cards. Yes, cable cards are a pain in the bahootie. Everything bad you've heard about them is (mostly) true. But, once you have them (mostly) working, the combination of cable cards and the TiVo is unbeatable. It gets you the whole HDTV experience plus everything that TiVo does that posers wish they could do.

For example, with two cable cards installed, TiVo can record two standard definition or HDTV programs at once, or record one program while you watch another. Some cable companies won't charge you extra for a standard digital cable box in addition to the two cable cards, so you even can have TiVo record two programs while you watch a third directly from the cable box. It's nice to have the cable box handy anyway, because you can't get pay-per-view directly from within TiVo.

The best part is that TiVo behaves like, well, TiVo. TiVo takes average reaction time into account when fast-forwarding or rewinding through a show. When you hit Play, it backs up to the spot you probably saw on the screen when you decided to hit the Play button. Add to that the season-pass approach to subscriptions and the intuitive interface—nothing even comes close to the user-friendliness of a TiVo.

Okay, so what's with this mostly stuff? I don't know if the cable cards are at fault or the cable company is at fault, but I've had two problems with cable cards. They tend to lose their authentication occasionally. They will re-authenticate themselves automatically, but your reception drops out in the meantime. I also do not receive a couple of premium channels I'm supposed to get, and the cable company hasn't figured out why. They're redundant channels, so it doesn't affect my viewing, but it's an annoyance.

The bottom line is if you want a real PVR, get TiVo. I've come around to GPLv3, and I'm in favor of the new license. But, it will be a sad day if GPLv3 forces TiVo to fork Linux in order to continue using it. TiVo has done a remarkable job with Linux and its hardware, and it should be rewarded, not punished.■

Nicholas Petreley is Editor in Chief of *Linux Journal* and a former programmer, teacher, analyst and consultant who has been working with and writing about Linux for more than ten years.

FANATICAL SUPPORT™

Joel Smith, CTO, AppRiver

"Luckily, the hurricane didn't blow us away. But Fanatical Support™ did."

"Hurricane Ivan hit our corporate offices in Pensacola, Florida and essentially shut us down. To our surprise, we got a call from Rackspace offering to handle our phone lines, our support requests, even our sales orders. Thanks to them, we were signing up customers as the hurricane was coming ashore. Now, if that's not fanatical, I don't know what is."

Going above and beyond to keep customers online is one definition of Fanatical Support. What will yours be?

Watch Joel's story at **www.rackspace.com/fanatical**
1-888-571-8976

rackspace® MANAGED HOSTING

Hear Yourself Think Again!

WhisperStation™

Cool... Fast... Silent!

For 64-bit HPC, Gaming and Graphic Design Applications

Originally designed for a group of power hungry, demanding engineers in the automotive industry, WhisperStation™ incorporates two dual core AMD Opteron™ or Intel® EM64T™ processors, ultra-quiet fans and power supplies, plus internal sound-proofing that produce a powerful, but silent, computational platform. The WhisperStation™ comes standard with 2 GB high speed memory, an NVIDIA e-GeForce or Quadro PCI Express graphics adapter, and 20" LCD display. It can be configured to your exact hardware specification with any Linux distribution. RAID is also available. WhisperStation™ will also make a system administrator very happy, when used as a master node for a Microway cluster! Visit www.microway.com for more technical information.

Experience the "Sound of Silence".
Call our technical sales team at 508-746-7341 and design your personalized WhisperStation™ today.

AMD 64 Opteron

intel inside XEON

Microway®
Technology you can count on℠