


MISTERHOUSE | F-SPOT | AJAX | KAFFEINE | ROBOTS | VIDEO CODING

LINUX JOURNAL™

Since 1994: The Original Magazine of the Linux Community

OCTOBER 2006 | ISSUE 150 | www.linuxjournal.com

An Excerpt
from Apress'
*Beginning
Ubuntu Linux:
From Novice
to Professional*



DIGITAL LIFESTYLE with Linux



- » F-Spot Tips » Working with Digital Images
- » H.264 Video Encoding for Low-Bitrate Video
- » Linux-Based Do-It-Yourself Robots
- » Share Music with Kaffeine, Amarok, Last.fm and more
- » Digital Convenience at Home with Open-Source Technology
- » Maddog's Travel Gadgets
- » Using MisterHouse for Home Automation

+ Doc Searls Breaks the **Marketing Matrix**





High-Performance and Enterprise Computing Under Your Control

Industry Leading 2P and 4P x86 Computing

Innovative server technology with outstanding performance and unrivaled memory scalability to accelerate compute and memory intensive applications.

2-way XtremeServer™ up to 64GB of memory



1U XtremeServer™

- Dual-Core AMD Opteron™ processors
- Up to 64GB of DDR2 533/667 memory
- Up to 1.0TB SATA or 292GB SAS
- 1 PCI-X and 1 PCI-Express x16
- Dual-port Gigabit NICs
- Hot-swappable drives
- ServerDome Management – IPMI 2.0
- Windows® or Linux OS

4-way XtremeServer™ up to 128GB of memory



3U XtremeServer™

- Dual-Core AMD Opteron™ processors
- Up to 128GB of DDR2 533/667 memory
- Up to 3.0TB SATA or 876GB SAS
- 3 PCI-X and 2 PCI Express x16
- Dual-port Gigabit NICs
- Redundant power supplies and fans
- Hot-swappable drives
- ServerDome Management – IPMI 2.0
- Windows® or Linux OS

4-way XtremeWorkstation™ up to 128GB of memory



XtremeWorkstation™

- Dual-Core AMD Opteron™ processors
- Up to 128GB of DDR2 533/667 memory
- Up to 2.0TB SATA or 584GB SAS
- 3 PCI-X and 2 PCI Express x16
- Two NVIDIA 4500/5500 graphics cards
- Redundant power supplies and fans
- Hot-swappable drives
- Windows® or Linux OS



Get the facts...

Performance Analysis and Value Proposition White Paper
Go to <http://www.appro.com>

AMD Opteron™ Processors: - AMD 64 Dual-Core Technology improve system throughput for faster networking connectivity
- Best performance per-watt, helping to reduce electricity costs while maximizing IT budget dollars



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



Today, Carlo restored a failed router in Miami, rebooted a Linux server in Tokyo, and remembered someone's very special day.

With Avocent centralized management solutions, the world can finally revolve around you. Avocent puts secure access and control right at your fingertips – from multi-platform servers to network routers, your local data center to branch offices. Our “agentless” out-of-band solution manages your physical and virtual connections (KVM, serial, integrated power, embedded service processors, IPMI and SoL) from a single console. You have guaranteed access to your critical hardware even when in-band methods fail. Let others roll crash carts to troubleshoot – with Avocent, trouble becomes a thing of the past, so you can focus on the present.

Visit www.avocent.com/special to download *Data Center Control: Guidelines to Achieve Centralized Management* white paper.



CONTENTS

OCTOBER 2006
Issue 150



COVER PHOTOGRAPHY: DIGITAL VIDEO CAMERA: ©ISTOCKPHOTO.COM/INGVALD KALDHUSSÆTER, LAPTOP: ©ISTOCKPHOTO.COM/ANASTASIYA MAKSYMENKO

FEATURES

50 CALL MISTERHOUSE TO REGULATE YOUR HEAT

One-Wire interfaces and MisterHouse come to the rescue for allergies.

David Lynch

58 DO-IT-YOURSELF ROBOTS WITH LINUX

Robotics made easy—especially if you get free sonar detectors.

Michael Surran

64 BUILDING A DIGITAL LIFESTYLE WITH OPEN-SOURCE TECHNOLOGY

Proof positive that nerds have all the cool toys.

Mike Diehl

68 MANAGE YOUR PHOTOS WITH F-SPOT

Get a taste of the power of Mono in this photo management tool.

Daniel Bartholomew

ON THE COVER

- An Excerpt from Apress' *Beginning Ubuntu Linux: From Novice to Professional*, p. 78
- F-Spot Tips, p. 68
- H.264 Video Encoding for Low-Bitrate Video, p. 74
- Linux-Based Do-It-Yourself Robots, p. 58
- Share Music with Kaffiene, Amarok, Last.fm and more, p. 32
- Digital Convenience at Home with Open-Source Technology, p. 64
- Maddog's Travel Gadgets, p. 22
- Using MisterHouse for Home Automation, p. 50
- Doc Searls Breaks the Marketing Matrix, p. 42

The competition doesn't stand a chance.



If you base deployment decisions on performance and price, Coyote Point's for you. We've cornered that market.

To prove it we asked The Tolly Group to evaluate our E350si application traffic manager against the competition. The results speak for themselves.

Throughput? Almost 40% more than others in our space. Cost of transactions per second? Up to four times less. Connection rate? In some cases, one-sixth the cost. One-sixth! And we're told Coyote Point is the #1 choice for today's open source networks.

But don't just take our word for it. Get the facts. Call 1.877.367.2696 or write info@coyotepoint.com for your free copy of the full Tolly Report.



CONTENTS

OCTOBER 2006

Issue 150

COLUMNS

24 REUVEN M. LERNER'S AT THE FORGE

JavaScript, Forms and Ajax

32 MARCEL GAGNÉ'S COOKING WITH LINUX

Your Media: Out of the Kitchen, and into Every Other Place



38 DAVE TAYLOR'S WORK THE SHELL

Analyzing Log Files

40 JON "MADDOG" HALL'S BEACHHEAD

What's in a Name?

42 DOC SEARLS' LINUX FOR SUITS

Breaking the Matrix

96 NICHOLAS PETRELEY'S /VAR/OPINION

MythTV Goes Amiss

IN EVERY ISSUE

12 LETTERS

16 UPFRONT

48 NEW PRODUCTS

81 ADVERTISERS INDEX

INDEPTH

74 ADVANCED VIDEO CODING ON LINUX

How do you make high-quality low-bitrate digital video?

Dave Berton

78 CHAPTER 10: PERSONALIZING UBUNTU: GETTING EVERYTHING JUST RIGHT

Whoever thinks you can't personalize a GNOME desktop needs to think again.

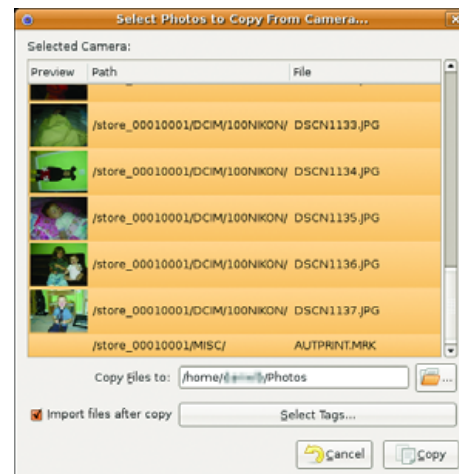
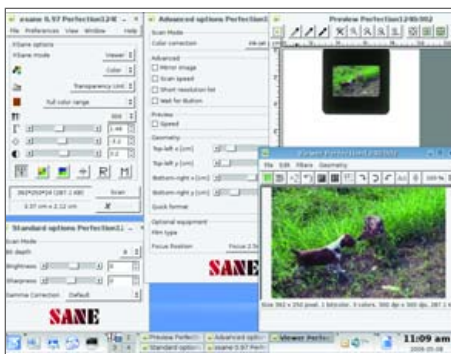
Keir Thomas



90 DIGITAL PHOTOGRAPHY AND LINUX

Whether you're Kooka or SANE, these digital photography tools could be for you.

Adrian Klaver



68 F-SPOT

Next Month

BUILDING DYNAMIC WEB SITES

Director of Web Technologies for Sun Microsystems Tim Bray talks with *Linux Journal* about XML, Java, Ruby, PHP, Perl, Web services and even his favorite distribution of Linux. Want to get more familiar with Ajax? Try implementing caller ID with Asterisk and Ajax. We also look into why Planetizen migrated to Drupal for content management. And, did you know that you can create Web sites with DocBook, XML and CSS?

There's more, as always. We look at how you can monitor your services and restart them automatically if any of them fail, and examine how Linux is affecting the telecommunications industry. You'll also find your favorite columnists stirring the pot and giving valuable tips.

USPS *LINUX JOURNAL* (ISSN 1075-3583) is published monthly by SSC Media Corp., 2825 NW Market Street #208, Seattle, WA 98107. Periodicals postage paid at Seattle, Washington and at additional mailing offices. Cover price is \$5 US. Subscription rate is \$25/year in the United States, \$32 in Canada and Mexico, \$62 elsewhere. POSTMASTER: Please send address changes to *Linux Journal*, PO Box 980985, Houston, TX 77098. Subscriptions start with the next issue.

From each flavor of Linux

to bite-size browsing on Nintendo DS™ -
Opera goes where you go.

Follow the standards.

Break the rules.

Test in Opera first.

Download the **free**
Opera browser -

www.opera.com



HOT WIRED FOR SERVICE



A Better Way
to Manage your Servers

Dedicated Servers
Starting at **\$39/mo.**

- No Setup Fees, No Contracts
- Free http and icmp monitoring
- Free incident response service
- Zervex ServerCP Online Control Panel Included



JTLnet
web hosting specialist

Real People Real Support™

-----Since 1998-----

www.jtl.net/lj
1-877-765-2300

LINUX JOURNAL

Editor in Chief

Nick Petreley, ljeditor@ssc.com

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

Contributing Editors

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

Proofreader Geri Gale

VP of Sales and Marketing Carlie Fairchild
carlie@ssc.com

Marketing Manager Rebecca Cassity
rebecca@ssc.com

International Market Analyst James Gray
jgray@ssc.com

Sales Coordinator Lana Newlander
ads@ssc.com

Regional Advertising Sales

NORTHERN USA: Joseph Krack, +1 866-423-7722 (toll-free)
EASTERN USA: Martin Seto, +1 416-907-6562
SOUTHERN USA: Laura Whiteman, +1 206-782-7733 x119
INTERNATIONAL: Annie Tiemann, +1 866-965-6646 (toll-free)

Advertising Inquiries ads@ssc.com

Publisher Phil Hughes
phil@ssc.com

Accountant Candy Beauchamp
acct@ssc.com

Linux Journal is published by, and is a registered trade name of, SSC Media Corp.

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

Subscriptions

E-MAIL: subs@ssc.com

URL: www.linuxjournal.com

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

USPS LINUX JOURNAL (ISSN 1075-3583) is published monthly by SSC Media Corp., 3262 Westheimer Rd, Suite 502, Houston, TX 77098. Periodicals postage paid at Seattle, Washington and at additional mailing offices. Cover price is \$5 US. Subscription rate is \$25/year in the United States, \$32 in Canada and Mexico, \$62 elsewhere. POSTMASTER: Please send address changes to *Linux Journal*, PO Box 980985, Houston, TX 77098. Subscriptions start with the next issue.

LINUX is a registered trademark of Linus Torvalds.



\$99



\$119



RouterBOARD 133 Trifecta

- 105mm x 105mm (4.1in x 4.1in)
- RouterOS Wireless AP/Firewall/Routing software already included
- Best price/performance ratio compared to similar products
- 175MHz MIPS CPU
- 3 miniPCI slots (one on bottom side)
- 3 10/100 Ethernet MDI-X ports
- 64MB NAND storage
- 32MB SDRAM
- Serial console port
- Passive PoE
- Input Voltage 9-30V
- Speaker, and Fan support

- Also available the RouterBOARD 133C with one ethernet, one miniPCI and 16MB SDRAM Perfect for CPE solutions. Only **\$69**

RouterBOARD 153

- 160mm x 160mm (6.3in x 6.3in)
- RouterOS Wireless AP/Firewall/Routing software already included
- 175MHz MIPS CPU
- 3 miniPCI
- 5 10/100 Ethernet MDI-X ports
- 64MB NAND storage
- 32MB SDRAM
- Speaker
- Serial port
- Passive PoE (803.3af power over datalines)
- Input Voltage 11-48V

RouterBOARD 44G

- a gigabit card with four independant ports
- up to 24 Gigabit Ethernet ports in a PC
- no more straight/cross cable problems, supports MDI/X
- server quality Realtek 8169SC chips
- an affordable Gigabit solution

www.routerboard.com
for extreme conditions

CIARA

DO MORE WITH LESS

For powerful servers that let you do more work on fewer systems, choose the **Dual-Core Intel® Xeon® Processor** in your **NEXXUS 4000®** Personal Cluster.

This Dual-Core Intel® Xeon® Processor based Personal Cluster server is the ideal platform for clients who have technical or mission-critical applications that require high processor count and fast interconnect. Based on a desk-side design, the NEXXUS 4000® is an alternative to traditional datacenter centric solution and the perfect system for IT or Research departments looking to leverage the benefits of Intel® standardization.



Visit us

For all your computer needs

Give us a call

CIARA-TECH.COM

1-866-789-7225

TECH

PERSONAL CLUSTER

8 PROCESSORS 16 CORES

\$19,995

8 Dual-Core Intel® Xeon® Processors 5148
32GB FBD ECC/Reg DDR2 533MHz
1TB (4 x 250GB) of Storage
Built-in Gigabit/KVM/USB Switch

Upgrade Your System

Additional 32GB Memory	\$6,599
Additional 1TB of Storage	\$1,099
InfiniBand Interconnect	\$3,455
(including PCI/Express HCA, Switch and Cables)	

- 150 Gigaflops¹ computing power.
- Up to 128GB² FBD DDR2 533/667MHz.
- Built in 8 or 16 Port Gigabit Switch.
(Optional 8 Port InfiniBand Switch)
- Up to 8 Serial ATAII HDD.
- Integrated KVM/USB Switch.
- Plugs into one single 110/220V outlet.
- Convenient desk-side form factor.



Core™, Core Inside, Centrino, Centrino Logo, Core Inside, Intel, Intel Logo, Intel Core, Intel Inside, Intel Inside Logo, Intel SpeedStep, Intel Vii, Itanium, Itanium Inside, Pentium, Pentium Inside, Xeon and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. (1) Important Information. All prices, specifications and promotional offers are subject to change without notice. Claim cannot be responsible for typography errors, photographs errors, pricing errors. All pricing in US dollar. Shipping and applicable taxes are not included. (2) Based on peak performance with Dual-Core Intel® Xeon® Processor 5148. Available with 4GB FBDIMMs.



Monarch Computer.com™



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



nero
7 ULTRA EDITION
with an upgrade to AMD Athlon™ 64!

Get Nero 7 Ultra, a Monarch shirt or hat, a Livestrong™ Wristband, and a FREE GIFT when you upgrade to Athlon 64/FX/X2/Opteron 100 (939) or any Socket AM2 bundle!

NEW! Enter to win 1 of 12 Autographed Lance Armstrong Jerseys!



Purchase any AMD Athlon™ 64/FX/X2/Opteron 100 (939) or Socket AM2 bundle from Monarch and get Nero 7 Ultra, a Monarch shirt or hat, choice of FREE GIFT and a LIVESTRONG™ wristband for helping Monarch support the battle against cancer. Monarch Computer and you can make a difference.

FREE INSTALLATION SETUP & TESTING BY CERTIFIED TECHS

FREE TECH SUPPORT!

Monarch makes it quick and easy to upgrade with FREE setup and testing on Motherboard Combos and \$45.00 build fee on Barebones.

AMD Motherboard Combos

FREE SETUP and TESTING

DFI NF4 Ultra Infinity w/ AMD Opteron™ processor 146 (939 pin)

Only \$225

Abit AN8-32X nForce4 w/AMD Opteron™ processor 180 (939 dual-core)

Only \$755



ASRock AM2NF4G-SATA2 with AMD Athlon™ 64 X2 processor 4200+ (Socket AM2)

Only \$259

Asus M2N32-SLI Deluxe w/AMD Athlon™ 64 processor FX-62 (Socket AM2)

Only \$1,075

Supermicro H8DAR-I w/ 2 x AMD Opteron™ processor 246 (940 dual-core)

Only \$649

Tyan S2892G3NR Thunder K8SE w/ 2 x AMD Opteron™ processor 270 (940 dual-core)

Only \$1,315

The AMD Opteron™ processor is designed for system stability and outstanding performance to help ensure that your IT environment continues to remain up and running.

AMD Barebone Systems

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



Antec Sonata II Piano Black Quiet ATX 2.0 Mid-Tower w/ 450W PS
Asus M2NPV-VM
AMD Athlon™ 64 3500+ processor (Socket AM2)
Starting @ \$344

Monarch Hornet Pro (Indigo) w/270W PS
ASRock AM2NF4G-SATA2
AMD Sempron™ 3000+ processor (Socket AM2)
Starting @ \$333

Antec Titan550 E-ATX Full-Tower/Server w/550W PS
Tyan S2877G2NR nVIDIA CrushK8-04
2 x AMD Opteron™ processor 250 (940)
Starting @ \$1,188

Tyan (Barebone) Transport GT24 w/500W PS
2 x AMD Opteron™ processor 280 (940 Dual-Core)
Starting @ \$2,791



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

Lowest Prices on AMD! www.monarchcomputer.com

AMD Athlon™ 64 X2/FX Dual Core Retail CPUs (939 and Socket AM2)
AMD Athlon™ 64 X2 3800+ (\$12K per core) \$152.00
AMD Athlon™ 64 X2 4200+ (\$12K per core) \$188.00
AMD Athlon™ 64 X2 4600+ (\$12K per core) \$240.00
AMD Athlon™ 64 X2 5000+ (\$12K per core) \$315.00
AMD Athlon™ 64 FX-62 (1MB per core) \$827.00

AMD Athlon™ 64/FX Single Core Retail Box CPUs (939 and Socket AM2)
AMD Athlon™ 64 3500+ (\$12K per core) \$91.00
AMD Athlon™ 64 3800+ (\$12K per core) \$111.00

AMD Opteron™ Dual Core OEM CPUs (939 pin)
AMD Opteron™ 165 1.8GHz \$249.00
AMD Opteron™ 170 2.0GHz \$361.00
AMD Opteron™ 175 2.2GHz \$492.00
AMD Opteron™ 180 2.4GHz \$659.00
AMD Opteron™ 185 2.6GHz \$840.00

AMD Opteron™ OEM CPUs (939)
AMD Opteron™ 144 1.8GHz (939) \$125.00
AMD Opteron™ 146 2.0GHz (939) \$157.00
AMD Opteron™ 148 2.2GHz (939) \$199.00
AMD Opteron™ 150 2.4GHz (939) \$304.00
AMD Opteron™ 152 2.6GHz (939) \$423.00
AMD Opteron™ 154 2.8GHz (939) \$623.00

AMD Opteron™ Dual Core for Socket F
AMD Opteron™ Model 2210 \$252.00
AMD Opteron™ Model 2212 \$373.00
AMD Opteron™ Model 2214 \$517.00
AMD Opteron™ Model 2216 \$691.00
AMD Opteron™ Model 2218 \$864.00
AMD Opteron™ Model 2220 SE \$1,153.00
AMD Opteron™ Model 8212 \$864.00
AMD Opteron™ Model 8214 \$1,153.00
AMD Opteron™ Model 8216 \$1,498.00
AMD Opteron™ Model 8218 \$2,127.00
AMD Opteron™ Model 8220 SE \$2,621.00

AMD Opteron™ OEM CPUs (940)
AMD Opteron™ 246 2.0GHz \$158.00
AMD Opteron™ 248 2.2GHz \$203.00
AMD Opteron™ 250 2.4GHz \$307.00
AMD Opteron™ 252 2.6GHz \$441.00
AMD Opteron™ 254 2.8GHz \$669.00
AMD Opteron™ 256 2.8GHz \$825.00
AMD Opteron™ 850 2.4GHz \$677.00
AMD Opteron™ 852 2.6GHz \$847.00
AMD Opteron™ 854 2.8GHz \$1,130.00
AMD Opteron™ 856 3.0GHz \$1,469.00

AMD Opteron™ Dual Core OEM CPUs
AMD Opteron™ 265 1.8GHz \$313.00
AMD Opteron™ 270 2.0GHz \$450.00
AMD Opteron™ 275 2.2GHz \$682.00
AMD Opteron™ 280 2.4GHz \$842.00
AMD Opteron™ 285 2.6GHz \$1,039.00
AMD Opteron™ 865 1.8GHz \$691.00
AMD Opteron™ 870 2.0GHz \$864.00
AMD Opteron™ 875 2.2GHz \$1,153.00
AMD Opteron™ 880 2.4GHz \$1,498.00
AMD Opteron™ 885 2.6GHz \$2,127.00

AMD Opteron™ Dual Core HE (Low-wattage) OEM CPUs
AMD Opteron™ 260 HE 1.6GHz \$470.00
AMD Opteron™ 265 HE 1.8GHz \$703.00
AMD Opteron™ 270 HE 2.0GHz \$862.00
AMD Opteron™ 275 HE 2.0GHz \$1,060.00
AMD Opteron™ 860 HE 1.6GHz \$864.00
AMD Opteron™ 865 HE 1.8GHz \$1,153.00
AMD Opteron™ 870 HE 2.0GHz \$1,498.00
AMD Opteron™ 875 HE 2.2GHz \$2,127.00

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

Components and Upgrades

1000s of In-Stock PC Parts and Accessories

101250 SilverStone Strider ST56F 560W ATX 2.0/EPS Crossfire Power Supply (Black)
\$129.55

150939 Western Digital Caviar SE 250 GB SATA3G 16MB Cache 7200 RPM (WD2500KS)
\$79.80

190156 Visiontek (ATI) Radeon X1300 PCI-E SFF 256mb GDDR2 DVI/TV-Out Retail
\$99.69

140123 512 MB DDR2 (533) PC2-4200 Corsair (VS512MB533D2)
\$51.60

101242 Lian-Li PC-G70B Aluminum Mid-Tower ATX (Black)
\$187.55

270270 Corsair 2 GB CMFUSB202GB Flash Voyager USB 2.0 Drive
\$58.62

150061 3ware 9590SE-8ML PCI-E x4 half-length 8-Ports
\$545.00

140803 1 GB (2 pcs 512) DDR (400) PC-3200 Corsair XMS (TwinX1024-3200C2P)
\$114.76


NEW! Section 508 Documentation Available for our PCs!

Educational and Government POS Welcome.

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



Commercial leasing available for purchases as low as \$1000.



Please note that prices are subject to constant fluctuation. Remember to consult www.monarchcomputer.com for the most current pricing. Pricing is based on the date of the issue in which the publication appears. Monarch Computer is not responsible for typographical or pricing errors. AMD, the AMD Arrow logo, AMD Athlon, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. All brands and product names are trademarks or registered trademarks of their respective companies. AMD model numbers indicate relative software performance among this AMD processor family. *Enhanced Virus Protection (EVP) is only enabled by certain operating systems including the current versions of Microsoft® Windows®, Linux®, Solaris®, and BSD Unix®. After properly installing the appropriate operating system, users must enable the protection of their applications and associated files from buffer overflow attacks. Consult your OS documentation for information on enabling EVP. Contact your application software vendor for information regarding use of the application in conjunction with EVP. AMD and its partners strongly recommend that users continue to use a third party anti-virus software as part of their security strategy. Please see www.monarchcomputer.com for important contest terms and conditions. No purchase required for entry.

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



Monarch's NEW EMPRO® 2 line of systems now features Next-Generation AMD Opteron™ processors with DDR2 and AMD Virtualization™!

Improving Direct Connect Architecture

- > For continued success in the enterprise
- Advancing Performance-per-Watt leadership**
- > Low-power, high-performing DDR2 memory
- > Consistent power roadmap with low-power options

Extending the Lead in x86 Virtualization

- > Founded on Direct Connect Architecture
- > AMD Virtualization is designed to improve business functionality and flexibility

Reducing Total Cost of Ownership (TCO)

- > One transition to your next stable platform
- > Seamless Dual-Core to Quad-Core upgrade in same thermal infrastructure
- > Improved Memory RAS and cost savings from DDR2



The AMD Opteron™ processor with Direct Connect Architecture provides industry-leading performance-per-watt and price/performance-per-watt.

"BOTTOM LINE: MUST BUY"

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



Jason Perlow
 Linux Magazine
 April 2005

Check out MonarchComputer.com's AMD Store for more information on Next-Generation AMD Opteron™ 1000/2000/8000 Series processors

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

You can be confident in knowing that AMD has longstanding partnerships with industry leaders -Microsoft and major Linux vendors.

**GET QUICK QUOTES
 Online or by phone:**

www.monarchcomputer.com

1-800-611-0875

Paypal - Visa - Mastercard - Discover - AMEX



LOOK FOR THIS SYMBOL
 ON LINUX COMPATIBLE
 MONARCH SYSTEMS!



**Monarch SuperMicro
 4-Way Custom 4U/Tower
 Server(PCI-E)**
 Ask for Part #: 80820



- SELECTED COMPONENTS:**
- SuperMicro AS-4040C-TRB 4U/Tower (Black)
 - 1000W Redundant Power Supply w/ PFC
 - Super H80CE Motherboard
 - Choice up to 4 AMD Opteron™ 800 Series Processors
 - Choice up to 64GB REG/ECC DDR Memory
 - Choice up to 5 SATA Hot-swap Drives
 - Dual Gigabit (Intel 82546GB) Ethernet Controller
 - nVidia® nForce Pro 2200 (CK804)
 - AMD8132 Chipset
 - Choice up to 2 Optical Drives
 - Choice of Floppy Drive
 - Choice of Sound Card
 - Choice of PCI-E Video Card
 - Choice of Linux or Microsoft OS
 - Choice of Network Options
 - Industry Standard Upgradable
 - Up to 3 year warranties available
 - 24/7 on-site service available

**Configuration Starting @
 \$3,298!**

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.

**Monarch Empro® 2
 Custom 2U
 Rackmount SATA
 Server** Ask for Part #: 80669



- AVAILABLE COMPONENTS:**
- Choice of 2U Rackmountable Chassis
 - Choice of Tyan or SuperMicro E-ATX Dual Opteron Motherboard
 - Up to 2 AMD Opteron™ 2000 Series Processors (Dual-Core)
 - Up to 32 GB DDR2 (400) PC2-3200
 - Up to 6.0 Terabytes of storage capacity when using SATA Hard Drives with options for RAID 0, 1, 5, 6 or 10 Setup
 - Up to 2 Optional RAID Controllers
 - Optional Optical Drives w/Software
 - Choice of Linux or Microsoft OS
 - Choice of Network Options
 - Optional Rail Kit
 - Industry Standard Upgradable
 - Up to 3 year warranties available
 - 24/7 on-site service available

**Custom configs
 Starting @ ONLY \$1,898!**

1 Configuration review

Our senior technical staff review every configuration to eliminate hardware and software incompatibilities.



2 Production scheduling

We allocate all parts for your system, and chart your system's assembly path through our production facility.



3 Board test

We assemble your motherboard, processor and memory and test these core components extensively. We also load the latest BIOS.



4 Build stage

All components for your system are assembled into your chassis. All cables are tied off and tucked away to increase airflow and cooling.



5 Burn in diagnostic

We combine hands-on diagnostics with a battery of automated burn-in testing to ensure all your components are operating properly together.



6 Software load

We load your OS onto your hard drive along with all factory tested updates and the most recent hardware drivers.



7 Quality control

Our QC experts put your system through a rigorous 62 point inspection to verify the system is in working order before shipping.



8 Packing & shipping

We expertly pack your system for secure and safe shipping using customized packaging and double boxing.



Size Matters

I am not writing to complain about the technical content of your magazine; it is excellent quality. I am not writing to complain about the new graphic layout of your magazine—layouts, “look and feel” all come and go. It neither adds nor detracts from the quality of the writing. I am not writing to complain about `/etc/rant`, or the views and opinions of Nicholas Petreley. Everyone is entitled to his or her opinion, and I have always respected Nicholas, even way back when he wrote for *InfoWorld*.

I am, however, writing to complain about the new size of the magazine. I somewhat agree with the opinion of Greg Bullough, who believed that his toilet tank was badly designed. However, toilet tanks have been around longer than *Linux Journal*, and even longer than Linux itself.

As an engineer, I am aware of the need to design products to operate in the environment they reside. Therefore, I plead with you to reduce the width of your magazine by about 1.25 inches (32 mm, or 0.0051 furlongs).

--

Henry Hojnacki

Keep Your Opinions to Yourself

I was reading this month's magazine, when I came across this new column: `/var/opinion`.

I just wanted to say that based on the editor's recent opinions in his column under the old name `/etc/rant`, there is a much better name for his column: `/home/petreley/opinion`. As the current name, this name is also LSB-compliant, but like his opinions should be, his home directory should be kept to himself (not public).

--

Guilherme DeSouza

Why Do You Talk in Parables?

I don't understand why Maddog has to talk in parables. All the other authors don't attempt to give their life story and speak in stories to make their point. This is a publication aimed towards a professional community. Get facts and write an article and not some senile story.

Also the `/var/rant` article is an okay thing. Editors are allowed to have opinions. Just as long as when it comes to editing the magazine they stay objective, which I feel that you have. But the Beachhead column has to go.

--

Michael S. Dunsavage

I and many others love Jon's style and look forward to the columns every month. See the next reader's advice.—Ed.

Read What You Like, Skip What You Don't

I've been reading this magazine for years and maintain a current subscription to it. Do I always agree with the content selections? No, of course not. Some issues center on something of no value to me, but the next always comes back centered on something that is more to my liking.

The secret to my success is if I don't like an article, I stop reading and turn to the next one. That's a pretty complex task; maybe I should force myself to read the article and write a flame mail to you guys like lots of the letters to editor.

As far as `/etc/rant`, sure the name might be off a bit, but it's the first article I read in every issue. Once again, I might not agree (hasn't happened yet), but it is your opinion after all. It must be hard for some people just to skip an article, even on the very last page, because it sure draws attention in the letters to editor section every month.

This is still the best Linux mag I see on the market, and I have no intention of canceling my subscription any time soon. Keep up the good work people.

--

Steve Bavalis

No Bad Karma Here

You are right, [Fedora] disk labels suck. Why Red Hat insists on them beats me.

I've installed [many versions of Red Hat] on many different hardware platforms, some of them many times. I've never had a Sendmail hang on boot issue. Could it be your bad karma?

--

dx

Bad Karma Strikes Others

Saw your column in *Linux Journal* and wanted to respond. I also have had problems over the years with Sendmail locking up about half the time when starting Fedora/Red Hat distros.

--

Christopher

Make It Simple, Stupid

You'd think [PC] hardware would be easy to understand and use. It isn't. No one wants to play well with others. It's like the PC world is living in never-neverland, and economics has nothing to do with [it].

No other consumer product makes these demands. I don't need to know how an engine in a car works and how to build one. I simply seat myself in the vehicle, fasten the seat belt, check the mirrors, put the key in the steering wheel lock and turn it to start the car.



» Enterprise Solutions Built on *OPEN STANDARDS*

» We can help you migrate to Open Source Solutions in no time. We have the Enterprise server and storage solutions you are looking for, at a price you never thought possible. By using open standard components and software, we can pass significant cost savings to you. And with our new state-of-the-art, fully automated production facility, we can build and **SHIP** hundreds of machines daily, further saving you money by not having to wait weeks, or even months for your mission critical systems to arrive. All solutions are built to your specifications, loaded with your choice of software, and backed by a industry leading 3 year warranty.

Open Source Storage is now Open Source Systems !

Come visit us at www.OpenSourceSystems.com



866-664-7867
Sales@OpenSourceSystems.com
www.OpenSourceSystems.com
1195 Borregas Avenue, Sunnyvale, CA 94089



I'd say the PC is a complicated dinosaur whose days are numbered. The economics of things, if nothing else, will force the PC to become a much simpler device to use, operate and own. There is a line being drawn in what one needs and what one can afford. As long as the brutes are determining the course of events, this will be the hard rule.

--

Wge

I have been predicting the success of computing appliances for years for these very reasons. Maybe someday my prediction will come true. There is hope. More on that in a future issue.—Ed.

Ruby Inspiration

As a contributing editor to *LJ* with an acknowledged fondness for Perl, I opened the July 2006 issue on Ruby and, Perl die-hard that I am, groaned. Despite whatever initial reservations I may have had, I did what I always do and read the magazine cover to cover. "Ruby as Enterprise Glue" by Maik Schmidt on page 58 clearly demonstrated just why Ruby is to be taken seriously, and I was so intrigued that I'm now working my way through the Pick-Axe book (*Programming Ruby* by the Pragmatic Programmers). Thank you, *Linux Journal*, for reminding me that the word "open" not only refers to my preferred technology, but to my mind, too.

--

Paul Barry

The Ultimate Do-It-Yourself Linux Box

I must say this was a nice article [August 2006]. However, there is one thing I'm missing in all Linux-related media so far: a test of server mainboards. Many server-bound extension cards (like the 3ware RAID controllers or those nice multiport Gigabit cards) are PCI-X, not PCIe. But, if you start looking for a mainboard suitable for a Linux system, you are completely on your own. Could you please give an overview of server-type mainboards that are suitable for Linux at some time? Something like "these boards work for most Linux distributions, a few boards are fully supported including built-in hardware watchdogs"?

--

Sven

Excellent idea, Sven.—Ed.

Ultimate Linux Laptop

I was disappointed to see that there was no Ultimate Linux Laptop [August 2006 Ultimate Linux Box issue]. We work hard to get most of the hardware working perfectly, including the Hot Keys, acpi events and hibernation, and provide a very simple way to get most multimedia to play easily, thanks to programs like MPlayer,

Xine and XMMS. Our goal is for customers to open their laptops and get to work. To me that is the Ultimate Linux Laptop.

I would like to make a few comments about your */var/opinion* column [August 2006]. We have evaluated many configurations and keep coming back to the Intel Centrino platform. Support for Intel Centrino is very good and is relatively easy to get the necessary components working. We have looked at AMD processors, and 64-bit capability might be nice, but what are the real honest reasons for running a 64-bit version on a laptop where maximum memory capacity is 2GB or 4GB? I can think of only a couple of good situations to run 64-bit Linux on a notebook.

To me, the the real promise of 64-bit computing lies in the ability to linearly access more than 4GB of RAM, and if notebook hardware isn't there yet, 64-bit notebooks really are not necessary at this time.

--

**Joe Ryner, R Cubed Technologies, Inc.,
Chief Technology Officer**

I think software is a bigger problem for 64-bit systems than memory, but I still agree. I run dual-core AMD64 machines for everything but my notebook, which is based on—you guessed it—Intel Centrino.—Ed.

Deconstructing Constructive Criticism

I just read your response to a letter about your Fedora rant. I fully agree that calling Fedora's approach to partition labels boneheaded is not the same as calling the developers boneheaded. However, I made the mistake of telling several police officers that chasing the crowd at a fourth of July celebration back toward an active lightning storm rather than let them walk "near" the fireworks drop zone that was obviously not going to be used unless and until the storm passed was "idiotic". I was nearly arrested for maligning a peace officer. The fact that I actually criticized a poor decision, not a person, one that was likely not made by these officers, was totally lost. As I did not want to spend the night of the Fourth of July in jail standing up for my right to not actually malign police officers, I beat a hasty retreat.

The distinction between constructive and destructive criticism seems to depend on whose ox gets gored. Most people read what they wish to hear or not hear with no connection to what is actually said or written, making it particularly dangerous to communicate with people with clubs and guns.

--

Dave Lynch

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 77253-3587 USA. Please remember to include your complete name and address when contacting us.

LETTERS TO THE EDITOR: We welcome your letters and encourage you to submit them to ljeditor@ssc.com or mail them to SSC Editorial, 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 +1 206-782-7733 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.

Linux laptops. Supported.

Supported:
Pre-configured Linux installation.
You choose your laptop.
You choose your distribution.
You customize your configuration.
Let EmperorLinux do the rest.

Supported:
Technical support
by phone and email,
manufacturer's warranty,
system-specific user's manual

Supported:
Port connectivity:
USB, PCMCIA, VGA,
Express Card, SVideo,
parallel, serial, FireWire

Supported:
True multiprocessing
with Intel Core Duo CPU,
up to 4 GB RAM

Supported:
Media cards:
Compact Flash,
Secure Digital

Supported:
Biometric fingerprint
(GDM login with PAM)

Supported:
X Windows at full
LCD resolution,
NVIDIA and ATI
3D acceleration,
OpenGL

Supported:
Power management:
suspend, hibernate,
processor control

Supported:
Connectivity:
Internal gigabit ethernet,
wireless a/b/g, Bluetooth,
EVDO mobile broadband

Supported:
One touch to control:
suspend, hibernate,
brightness, volume,
external VGA, wireless

Supported:
Internal optical drive:
CDRW, DVD±RW

Supported:
Laptops from top tier
manufacturers you
know and trust:
Dell, Lenovo, Sharp,
Panasonic, Sony

Since 1999, EmperorLinux has provided pre-installed Linux laptop solutions to universities, corporations, and individual Linux enthusiasts. We specialize in the installation of the Linux operating system on a wide range of the finest laptops and notebooks made by IBM, Lenovo, Dell, Sharp, Sony, and Panasonic. We offer a range of the latest Linux distributions, as well as Windows dual boot options. We customize each Linux distribution to the particular machine it will run upon and provide support for: ethernet, wireless, EVDO mobile broadband, PCMCIA, USB, FireWire, X-server, CD/DVD/CDRW, sound, power management, and more. All our systems come with one year of Linux technical support by both phone and email, and full manufacturers' warranties apply.

Visit www.EmperorLinux.com or call 1-888-651-6686 for details.

EmperorLinux

...where Linux & laptops converge

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

Model specifications and availability may vary.



diff -u

WHAT'S NEW IN KERNEL DEVELOPMENT

Just as ext3 emerged as the successor to ext2 after **Linus Torvalds** said he didn't want additional features like journaling to clutter up a stable and functional filesystem, it now looks like **ext4** will emerge as the successor to ext3. This goes directly counter to the desires of the ext3 maintainers, who don't want to give up their user base or have to struggle with maintaining divergent source trees between two projects. But, Linus doesn't consider these to be serious objections, and he is unwilling to sacrifice the stability and reliability of ext3 for the convenience of its maintainers. Assuming the controversy doesn't tip the other way, we can look forward to an ext4 that can hold 1,024 petabytes, instead of ext3's paltry 8 terabytes, and that supports filesystem extents, a technique of preventing disk fragmentation and slow-down.

Intel has launched a project to support **WiMedia Ultra Wide Band (UWB)** and the wireless USB standards. UWB is a wireless technology with a very short range, useful primarily for communicating within a single room. The hardware supporting this technology is so new that it is apparently still hard to find. But, Intel is gearing up to support it and has invited the Linux community to contribute to its effort.

Theodore Y. T'so has been working on removing as much data as possible from disk **inodes**, and keeping them elsewhere in the kernel. The idea behind this is that files use a lot of inodes, and accessing a lot of files uses a lot of RAM for the inodes. If inodes were smaller, the system would need much less RAM to deal with open files. The best part is because all filesystems use inodes, shrinking the inode data structure will improve the memory usage of every filesystem out there. A bunch of folks, including **Alexander Viro**, have been helping to trim the inode data structure. Linus Torvalds also apparently supports the work, and Ted has very quickly submitted some patches to accomplish a lot of what he originally set out to do.

Sean Estabrooks has written a tool to import **Perforce repositories** into **git** and has posted his code and documentation for inclusion into the git tree.

Jon Smirl has begun the arduous task of attempting to import eight years of **Mozilla** development into git from a CVS repository. Whether the resulting repository will be used for ongoing Mozilla development remains to be seen.

Paul Mackerras has enhanced **gitk** to show the nearest tags (typically official version numbers) to a given commit. This allows users to identify more easily which release was first to accept a particular patch. Because this feature requires a significant amount of computation, Paul has set it to run in the background and update the display once the information is derived.

It's now possible to **alias git commands** (with all their command-line arguments) to other names, effectively creating customized commands. This ability to create short, quick alternatives to commonly used or favorite commands can save a lot of time if, like Linus Torvalds, you have to process massive quantities of patches on a daily basis.

There's been a small amount of confusion over the way git should interpret the **changelog** text portion of a patch. **Eric W. Biederman** recently patched git to interpret From headers anywhere within the changelog entry, as indicating patch authorship. But, Linus has insisted that the From header should indicate authorship only when it appears at the very top of the patch. The reason for this, he says, is that git (and really any version control system) should never guess at the meaning of data. As he puts it, "SCMs are not about guessing. They are about saving the *exact* state that the user asked for. No 'let's try to be nice', no gray areas."

The **diff output** of git can now be colored, though the precise command-line option to do this is still being hashed out. And, as Linus said, the initial color-set selected for git's diff output "will make most people decide to pick out their eyes with a fondue fork". So, clearly there is some tweaking still to be done.

—Zack Brown

SO DESIGN SOME PENGUIN ONES, SOMEBODY

John Fluevog (www.fluevog.com) would rather have a tongue in cheek than foot in mouth. Else how to explain the shoemaker's felicitous approach to his business, his shoes and his approaches to making and selling them.

The About page on Fluevog's site begins:

John and his longtime friend and ex-partner, Peter Fox, go back to biblical times. In those days, such scriptural characters as Methuselah lived to be, quite literally, thousands of years old. John, himself, is thousands of years old and Peter Fox is, in fact, Methuselah. Before they founded Fox & Fluevog, in the ancient biblical age of 1970, they were shepherds. Actually, they worked at Sheppard's, a venerable Vancouver shoe emporium, albeit a tad on the conventional side.

The story goes from here to Fox & Fluevog, "the very coolest place to treat the feet and has remained so ever since". Then:

Around 1980, Peter heard the siren song of New York City and moved there to open his own store, specializing in design of Ladies' (particularly wedding) shoes. John took over the business and built it into the multinational mega corporation it is today, along the way forging his undaunted reputation for the world's most distinctive shoes.

The FAQ below has questions such as "What's the

capital of Uruguay?" (The answer is in Spanish.) And, "Are John Fluevog Angels really Satan resistant?" (Answer: "You're wearing Angels right now, right? Look around—do you see him?")

Among the many delightful pages on the extensive Fluevog.com site is one (www.fluevog.com/files_2/os-1.html) that goes:

You've Heard of Open Source Software? Get Ready for John Fluevog's Open Source Footwear.

"But how can a software concept work with shoes?", we hear you gasp in astonishment. Keep reading.

YOU ALL KNOW how open-source software works, right. *Sure you do*: everybody has access to the program's source code and anybody who thinks of an improvement can just write it up and send it in. If others use it, Yahtzee! Everybody wins! It's simple, it's people-driven, it's non-monetary and it makes software people really want. *We're going to make shoes the same way.*

In a word balloon from a video frame of John Fluevog speaking:

It's a concept that's well, maniacally tremendous: got an idea for a shoe? Even just for part of a shoe? Scribble it down and send it to me. I don't care if it's on a bar napkin, as

long as I can make it out.

ARE YOU FRUSTRATED, not finding the shoes you really want? Is your imagination ahead of the whole shoe industry and you're sick of waiting for them to catch up? Here's your chance to go over their heads and deal with someone who really cares what you want. All you need is a brilliant idea. Fax it, mail it, upload it, e-mail it, bring it in—just get it to me.

These lead to an updated Results page where readers can vote on submitted designs.

And below that is a link to a page that explains open source as "both a software philosophy and an important progressive movement".


Not coincidentally, the Fluevog site runs fast for a graphics-heavy one. It should be no surprise that it's served by Apache off Linux servers in Vancouver, BC. Same with Peter Fox Shoes.

By the way, if your shoes really do get produced, John names them after you, and the design goes into the public domain—as well as admission to what John modestly calls "the hallowed ranks of Fluevog Design Alumni".

If any of you submit designs, and they turn into real Vogs on real feet, send photos to ljeditor@ssc.com, and we'll show them here.

To submit your order, visit this page: www.fluevog.com/files_2/os-1.html.

—Doc Searls



What could YOU do with **16** cores



THINKMATE 8-WAY WORKSTATION FEATURING AMD OPTERON™ PROCESSORS

This system offers nothing less than the flexibility and power to meet or exceed your computing requirements. Need a 5U rackmount? Need a full-scale tower? It's convertible to both form factors!



Better Efficiency, Greater Productivity, and Enhanced Scalability

Operating System
Microsoft Windows, Red Hat Enterprise Linux, and Sun Solaris Operating System Configurations

System Memory
Supports up to 128GB PC2700 DDR ECC Registered Memory

Video
PCI Express x16 or Integrated ATI Rage XL 8MB

Storage
Up to 8 Removable Serial ATA or SCSI Hard Drives

Power Supply
1350W 3+1 Redundant Power Supplies

Thinkmate Warranty
Thinkmate systems are warranted against defects in materials and workmanship for three years following the date of purchase.

- 3-Year Advanced Replacement of Defective Components
- 3-Year Technical Support
- Optional 3-Year Onsite with IBM Global Services
- Optional 4 Hour Response

- 8P/16C system with Dual-Core AMD Opteron™ processors Model 885 2.6GHz 2x1MB Cache
- 128GB (32 x 4GB) PC2700 DDR ECC Reg. Memory
- PNY nVidia Quadro FX 5500 SDI 1GB Graphics Adapter with Next-Generation Vertex and Pixel Programmability
- 8 x 300GB Seagate Ultra320 SCSI 10,000 RPM Hard Drives with Adaptec 2230SLP Ultra320 SCSI RAID Controller Card

**16
Cores**

\$62,999

- 8P/16C system with Dual-Core AMD Opteron™ processors Model 885 1.8GHz 2x1MB Cache
- 8GB (16 x 512MB) PC3200 DDR ECC Reg. Memory
- PNY nVidia Quadro FX 4500 512MB Graphics Adapter with High Precision Dynamic Range Imaging (HDPR) Technology
- 4 x 400GB Seagate Serial ATA 7,200 RPM Hard Drives with 3ware 9550SX-8LP Serial ATA RAID Controller Card

**16
Cores**

\$17,999

THINKMATE

(800) 371-1212
WWW.THINKMATE.COM

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

LJ Index, October 2006

1. Linux-based services among the ten most reliable hosting companies: **5**
2. BSD-based services among the ten most reliable hosting companies: **4**
3. Windows-based services among the ten most reliable hosting companies: **1**
4. Linux-based services among the top 50 hosting companies with the fewest failures: **21**
5. Windows-based services among the top 50 hosting companies with the fewest failures: **13**
6. BSD-based services among the top 50 hosting companies with the fewest failures: **7**
7. BSD-based services among the top 50 hosting companies with the longest uptimes: **43**
8. Linux-based services among the top 50 most requested sites: **29**
9. Apache's percentage of server market share: **73.24**
10. Billions of dollars in Linux server sales for Q4 2005: **1.6**
11. Billions of dollars in Linux server sales for all of 2005: **5.7**
12. Position of Kambja, Tartumaa, Estonia among the most "Penguin-Friendly Places" on Earth: **1**
13. Position of Antarctica among countries with the highest densities of Linux users: **1**
14. Highest uptime hours for machines followed by Linux Counter: **3,802.2**
15. Years on the highest uptime machine: **10.4**
16. Firefox + Mozilla browser share percentage of those visiting W3Schools.com in June 2006: **24.9**
17. Linux OS percentage of those visiting W3Schools.com in June 2006: **4.4**
18. Mac OS percentage of othose visiting W3Schools.com in June 2006: **3.6**
19. Percentage of LinuxDevices survey respondents that have used Linux in embedded projects or products: **47**
20. Percentage growth in Linux embedded use over last year's survey: **2**

Sources: 1–8: Netcraft.com (as of June 14, 2006) | 9–11: ServerWatch.com | 12–15: Linux Counter (as of July 10, 2006) | 16–18: W3Schools.com | 19, 20: LinuxDevices

—Doc Searls

They Said It

The people are the network.
—Erik Cecil, arch-econ e-mail list

Every night there's a Mashup get together, or a TechCrunch party, or it's Tag Tuesday, or SuperHappyDevHouse, or SXSW, or this conference or that conference. And this stuff is fun. It's a real community. But all of these things are great by themselves, but terrible in combination. I see some entrepreneurs in photos from every *single event*. Who's talking to the users, writing the code, tweaking and retweaking the UI? It ain't the Chief Party Officer.

—Caterina Fake,
deathsxpirl.blogspot.com/2006/04/
entrepreneur-engineer-connector.html

There has grown up in the minds of certain groups in this country the notion that because a man or a corporation has made a profit out of the public for a number of years, the government and the courts are charged with the duty of guaranteeing such profit in the future, even in the face of changing circumstances and contrary public interest. This strange doctrine is not supported by statute nor common law. Neither individuals nor corporations have any right to come into court and ask that the clock of history be stopped, or turned back, for their private benefit.

—Robert Heinlein



Are you
shocked
by the
high cost
of iSCSI &
Fibre Channel
storage?

AoE is your answer!

ATA-over-Ethernet = simple, low cost, expandable storage.

www.coraid.com



EtherDrive® SR1520

- RAID enabled 3U appliance with 15 slots for hot swap SATA disks
- Check out our other Storage Appliances and NAS Gateway



1. 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
for more information.



1.706.548.7200

The Linux Storage People

www.coraid.com

Bill's Retirement

Dear Bill,

I read in the paper the other day that you are retiring. Good for you! I have not quite gotten to that age yet, but I envy those who have put a little away so they can take early retirement.

Now, you know that you and I have not seen eye to eye on how to make money over the years. I believe in making people pay money only once, if that, and you seem to believe that they should pay you again and again for the same products. Well, that is water over the dam now, and I just want to make sure that you get on the right path to retirement.

The first thing I should talk to you about is saving money. Now that you are retired, that income will not just be rolling in as before. You should think about how to save money. I recommend looking closely at this "Free Software" thing. That is where you pull down the bulk of the software off the Internet (you know, that thing you said would never catch on) and then decide if there are any small changes or improvements you need to make to the software. If there are no changes, you have it for free! If there are some small changes, all they ask is that you give those changes back to the group. Nothing wrong with that, is there Bill?

Another thing you should think about is diversifying some of your investments. This is what my stock advisers keep telling me all the time. They tell me that when I have a lot of income coming in, it is okay to play around with those high-risk, high-return stocks, but when you retire, you should think about that solid, steady income.

I know that a lot of your money is tied up in that company you started. That is all well and good, but I have experienced that sometimes things go downhill after you leave, and you should not put all your eggs in one basket. I have heard, for instance, that there is this really interesting startup you might want to put a little seed money into, called Google. Lots of people are talking about Google, and I think it could be a good investment for just a little of your "mad money".

I have also heard that you are going to do a lot of volunteer work with charity. Now those charity people do a lot of things in foreign countries, and believe it or not, those countries do not speak English all the time. Once again, you should take a look at that Free Software stuff, as it allows those countries to change the software to meet their

own cultural needs. And, if you teach them how to do it themselves, they can save even more money. "Teach the man to fish", Bill, remember that line?

This is another reason for you to divest some of your investments, Bill. You have too much tied up in certain companies that may cloud your thinking. With as much money tied up in that company you started, you might think that the best thing for these educational needs, health needs and job needs are to give them copies of the products produced by your old company. You need to clear your head! You need to think outside the box! You need to embrace the idea that Free Software is the way to go for these charities!

Just think for a moment what happens when you give some of your products to a poor student in the Congo, and then the WPA program starts acting up on him. Who is he going to call? And without the source code to WPA, how is he going to use that software you gave him? Surely you can't expect him to pay the license fee for the update?

I found some students in South Africa that were studying photography. When they were using the software in their school, everything was fine. But when they went home to practice on their own computers, they had to use pirated software! So I gave them all copies of GIMP, the free software, and they were pirates no more! See how simple that was?

Now I heard that your good friend Mr Buffett gave somewhere around \$35 billion to your charity fund! That was awful nice of him. I bet that he has no restrictions on his gift, like making any of the charitable organizations or groups he is helping buy any of his CDs, or go

to any of his concerts. You need to match his generosity by making your donations without strings also (Get the joke? Guitar strings?).

So diverse yourself of the temptation, and invest in Free Software companies like Red Hat Software or even Novell. This will make your charitable donation stretch the farthest.

Speaking of charitable donations, I wonder how much money the Free Software people would have generated if they had sold their software again and again like you did instead of contributing it? Do you think it might have topped your donations? Mr Buffett's? A lot of margaritas there! So much about your donations in the press and so little about theirs! But I bet you can fix that when you start supporting Free Software.

Well, that is about all the news and advice I have to send you for right now. I will note that the papers said it would take you two years to remove yourself from the company you created. My gosh! I am glad you have never been hit by a truck! The company might have collapsed without you! If this is true, it is probably about time for you to move on. But I would even more strongly suggest that you diversify your portfolio.

Say hello to Melinda for me, and tell me how she liked the penguin earrings I sent to her.


Warmest regards,
maddog

P.S. My secretary tells me that it was Warren Buffett that gave you the donation, and not Jimmy. Too bad. I never met Warren personally, nor played cards with him, but I think Jimmy would have been more fun at your retirement party. And, you can forget the joke about the guitar strings.


USER FRIENDLY by J.D. "Illid" Frazer



LINUX JOURNAL EDITION



The Industry Leader for Server Appliances



Custom server appliances or off the shelf reference platforms,
built with your image and software starting under \$1,000.

From design to deployment, we handle it all.

Delivering an appliance requires the right partner. MBX Systems is the right partner. We understand that by putting your name on our hardware, you're putting your reputation in our hands. We take that seriously. We provide the services you need to support your customers. Better than the competition. You never even

need to touch the hardware. Engineering. Design. Deployment. We handle it all, so you can focus on what's important to you. Your software. Your sales. Your success.

Visit us at www.mbx.com or call 1-800-681-0016 today.



www.mbx.com | 1-800-681-0016

What's in My Backpack

Having traveled to more than 60 countries during the past ten years, and most of them more than once, I have gathered quite a few things in my backpack that allow me to leave the house feeling confident I can handle whatever comes my way electronically. Planning ahead helps, and things have changed over the past ten years, but I noticed lately that I have not added or left out anything from my backpack for about a year now, so things seem to be fairly stable. Keeping these items in the backpack means I have them ready at all times for that “next trip”.

The first thing is the backpack itself. After trying more than a few of them and having them fall apart under heavy use, I finally settled on the Wenger SwissGear Synergy Model #GA-7305-14. I do not know if it was really made for the Swiss Army, but it sure seems that way. With lots of pockets internal and external for stuff, it holds cables and various wires in their own places so I know where they are. It also has really strong straps and a carrying handle that is wrapped around a steel cable—talk about solid! The zippers are big-toothed and rugged, good for squeezing the last bit inside the bag.

The model I have seems to be missing from the Wenger site, but I found it at www.avenuesusa.com/wenger/7305.html, and they are carried at a lot of different retailers, such as Circuit City, Best Buy, Staples and so on.

One warning though, the main pocket is said to fit “most 15–15.4-inch screens”, and some of the on-line reviews talk about the “snugginess” of their 15” and even 14” notebook, so if you have a notebook with a very large screen, this may not be the backpack for you. I do not go for the largest screens on a laptop, because I expect that the person in front of me on an airplane will recline his or her seat and crush my laptop screen. My 12.1” laptop screen avoids this.

The backpack has a separate pocket in the top that carries an audio player, a set of earphones and the charger for the player, with a hole in the backpack to thread the earplugs through. I have a set of permanently installed earphones, so I can listen to music as I walk along and not have to unthread them when I get to where I am going. If I leave the backpack some place, I just unplug the earphones, grab the player and another set of sound-reducing earphones from the pocket and go where I need to travel.

The audio player itself is an iRiver H340, with a 40GB drive. It is one sweet unit, having the ability to play WMA, MP3 and Ogg formats, a built-in radio and a microphone to record. Unfortunately, it records only in MP3 format, but someday I hope iRiver will give its customers the chance to choose their recording format also. I ripped my 400+ CD collection into Ogg format and loaded it onto the disk. It is only about half full, so I can listen to music for more than 150 days continuously without having to hear the same song by the same artist. And, I also can listen to radio, of course. I chose iRiver mostly due to the Ogg support. I wanted a format that was truly free, and for which I could get all the software I needed to encode and decode it in source code format, and without patent issues. Unfortunately, the iRiver people do not make the H340 anymore, but they do have additional models at www.iriveramerica.com that support Ogg. You also can find a lot of Ogg players at wiki.xiph.org/index.php/PortablePlayers.

I also carry a Belkin FM Tunecast transmitter (Model F8V3080), which allows me to broadcast my iRiver (or laptop, or anything with an 1/8-inch stereo jack) to an FM radio, so I can easily listen to *my* music in the hotel room or automobile. This unit is only for sale and use in the United States, but perhaps other countries have similar

units. Finally, I keep a stereo splitter, with 1/8” stereo jacks in the same pocket, along with a cheap set of earbuds that are very light. You never know when a friend may want to listen to your stereo, and the idea of sharing one set of stereo earbuds between two people does not cut it with someone who wants to hear *both* Simon and Garfunkel at the same time.

In the front pocket of the backpack, I keep two Kensington cable locks, combination type—one lock for the notebook and one for the media slice that comes with the notebook. I can run one or both of the cables through the steel cable of the backpack handle and secure that at the same time. I chose a combination lock cable because of my usual habit of losing keys. You can find these at us.kensington.com/html/11209.html, but other manufacturers make them too.

The next pocket has a set of retractable cables from Ultra. This is a 13-piece set that comes in its own carrying case. It includes a stereo headset with microphone (good for listening and recording voice for my notebook and also good for VoIP), an RJ-11 cable, RJ-45 cable (with separate crossover converter), a FireWire Cable, two USB cables and a bunch of adapters for various sizes of USB and FireWire. The collection is great, I have never needed any other “connecting cables”, and the holder keeps them all organized and ready to go (I often lose cables too). You can find it at www.ultraproducts.com/product_details.php?cPath=9&pPath=148&productID=148 for \$39.99 US, but I have noticed that in the *SkyMall* airline shopping magazine, they are listed for a lot less, and for those of you who do not fly much, here is where you can find them on-line: www.skymall.com/webapp/skystore?process=prodDisplay&action=&pid=69697606&catid=14101.

The next thing is pure “maddog”, but I told you that I would tell you everything I keep in the backpack....a roll-up piano from Hecsan. This is my feeble attempt at trying to learn another keyboard, the musical keyboard. Don't ask me to play something for you, the time that I have had to practice has been nil, but I keep it there “just in case” I get the time. You can see them at www.hecsaninc.com/home.html. I have the original model, purchased several years ago, and although it is good, I can see that they have improved it as well.

Moving closer to my back in the backpack is a larger pocket. Here I keep most of the “goodies” except the notebook itself, which I keep in a separate pocket in the backpack.

The first thing out of the pocket when I travel is a Kensington adapter plug. Now, there are a lot of adapters on the market, but they usually come in many pieces. You know how I am with keys and cables. This plug is all-in-one and supports the plugs for 150 countries (there are approximately 200 countries in the world), most exceptionally South Africa, which has the strangest plug I have ever seen. I have the older model (model number 33117) that retails for \$19.99 US, but Kensington also has another model adapter so you can plug in your USB chargeable device and charge it without having to turn on your notebook (model number 33346), which retails for \$10 US more. Both can be seen at us.kensington.com/html/11172.html.

Into the Kensington, I plug a three-outlet power strip that is very small and light, which handles 110 up to 240 volts, but uses US plugs. Then I plug in all my US-plugged devices to this. Now you can see why I need only one adapter plug. When I am going to plug in only my notebook, I take out just the Kensington adapter and do not have to bring out the power strip. Of course, if you are from a different country, you can look for a power strip that supports your type of plug, but make sure it can handle all the voltages.

APC (www.apc.com) also makes an interesting UniversalPlug, but

do not throw away the little picture of how to put it together for all the different socket types of the world. If you do, you will need to be sure that you have a friend that is good at Rubik's Cubes, in order to get the plug into the proper configuration for any socket. I try to keep all three pieces (the plugs and the instructions) in one pocket of the backpack. The URL is www.apc.com/resource/include/techspec_index.cfm?base_sku=INPA&tab=models.

Often I get to a place where there is broadband Internet, but not wireless broadband, only cable. How quaint. So I bring along my D-Link Model DWL G730AP wireless access point, router and client. I can set up a wireless bubble for myself and whoever else wishes to use it. This is also useful in a hotel room that I might be sharing with a group of people—that might be interpreted wrong—I meant a suite of rooms that I am sharing with a group of people. You can find this at www.dlink.com/products/?sec=0&pid=346 and of course other manufacturers make similar devices.

When I am looking for a hot spot, I can use my Canary Wireless HS10 Hot Spot finder (www.canarywireless.com). It has the capability of finding the hot spot, showing whether it is "open" (there is that word again) or "secure", and the signal strength. By continuing to push the power button, it also shows all of the hot spots in that area, not just one, and it gives you the ESSID and channel of the hot spot. Some people simply use their notebook and appropriate software to do this, but the Canary tells me if it is at least worth hauling out the notebook. Of course, my notebook does B/G and A, and Canary does only B and G, but most hot spots are only B and G anyway. One problem with the Canary is that the power button is not recessed enough, so sometimes it gets pushed inside of the backpack and uses up its battery. Careful positioning of the Canary as I put it back keeps this from happening.

I also carry a Delorme Earthmate GPS unit (USB), so I can run gpsdrive to find out where the heck I am sometimes. And, I carry an old IBM VGA camera (Model XVP610) for a portable Webcam. It was very cheap when I bought it, and if it breaks I will not get upset.

Finally, I normally carry in my backpack a small case for my digital camera, and in that case I also carry a small recharger for AA batteries and a couple of sets of those batteries. Most of the portable devices that use batteries seem to use AA or AAA, and the recharging unit handles both, four at a time. A PCMCIA card makes the transfer of pictures to my notebook very quick.

And, there are also a few other dongles for recharging devices, such as cell phones and PDA devices. I have not looked at the iGO power devices.

The notebook itself is an IBM X31 with a "media slice" that provides extra ports, an extra battery and a removable device bay for CD/DVD reader/writers. I bought it almost two years ago, and knowing that it would be my main system for a while, I bought top of the line, with 2GB of main memory, an 80GB disk, 802.11 b/g/a wireless networking, FireWire, USB 2.0 and IR. I also prefer the ThinkStick of the IBM over the glidepad for "mousing", but when I am some place where there is a bit more room, I also have a three-button optical mouse (Newpoint, www.newpoint.com/catalog.jspa?webcat=3&itemNumber=230828&method=showProduct) that has a self-retractable cord and a clear case so you can see the components (geeky, I know). The X31 also has a ThinkLight built in, which illuminates the keyboard for typing at night when you just can't "remember" where that particular key is.

I chose the IBM because of its titanium case, good support for Linux and good field support. Recently, I had to replace the keyboard because the oils in my hand dissolve plastic ("a mutant, you are maddog!"), and they shipped me the field replaceable keyboard overnight. The power supply for the laptop is not only good for the 110–240 volt crowd, but also has an adapter for the car cigarette lighter socket and the airplane seat socket, so I do not have to have a separate power adaptor to use it in the car or airplane. This is what every vendor should do for its laptop customers.

I do not do floppy drives anymore, opting for a USB stick, which people are now giving away from time to time, but taking off the media slice allows for the X31 to become very small and very light. If I need a floppy drive, there are plenty that are USB-based and work fine with Linux.

So there you have it, a combination of components that fit in my backpack. The next time you see me running through the airport, you will know that it is not because I forgot a needed piece of equipment!

—Jon "maddog" Hall



axigen

AXIGEN is a fast, modular, production-class Mail Server including:

- (E)SMTP, POP3, IMAP4, WEBMAIL
- LIST SERVER
- ADVANCED FILTERING SYSTEM
- SSL/TLS ENCRYPTION
- WEB / CLI ADMINISTRATION



AXIGEN MAIL SERVER

TRIAL VERSION

- FULLY FEATURED
- UNLIMITED USERS
- SUPPORT INCLUDED

www.axigen.com

Visit now to start your free trial!



REUVEN M. LERNER

JavaScript, Forms and Ajax

Dynamic user detection and registration made easy via Ajax.

Last month, we began to dip our toes into the water of Ajax, the shorthand name for asynchronous JavaScript and XML, which has taken the world of Web development by storm. Ajax applications are Web applications in every way, depending on the underlying combination of HTML, HTTP, URLs, JavaScript and CSS that provide a modern Web infrastructure. But, they also rely upon several of the features of modern JavaScript, including its ability to rewrite Web pages and also to make HTTP requests asynchronously.

And, indeed, this asynchronous behavior is what makes JavaScript—and Ajax, for that matter—so exciting for Web developers. No longer are we stuck with the modern equivalent of old-style 3270 terminals, with execution taking place on the server only when we move from one page to another. Now a Web page can be updated without having to reload the page.

Ajax is not a technical revolution, so much as a conceptual one, bringing with it new expectations from users and paradigms for developers. All of the technology behind Ajax has existed for several years, but it is only now that we are starting to take advantage of it in Web applications.

This month, we start to look at one simple example of an Ajax application, in a context that will probably be familiar to most Web developers: asking users to register with a Web site. By the end of this month's column, you will see how we can combine server-side programs, an HTML form and JavaScript to check the validity of a form before it is submitted. Next month, we will see how we can use Ajax to overcome the fatal flaws associated with this implementation, improving the efficiency, robustness and security of the application all at once.

Registering Users

If you have been developing Web sites for any length of time, you probably have needed to create a login system. Such systems come in all shapes and sizes, and require different levels of security. For our example application, we assume that each user has a user name and password. Actually, as you will see,

we don't care that much about the password; the key issue here is the user name, which must be unique.

The first thing to do is create a simple table in our database to keep track of users:

```
CREATE TABLE Users (
  id SERIAL NOT NULL,
  username TEXT NOT NULL CHECK (username <> ''),
  password TEXT NOT NULL CHECK (password <> ''),
  email_address TEXT NOT NULL CHECK (email_address <> ''),

  PRIMARY KEY(id),
  UNIQUE(username)
);
```

The above table, defined using PostgreSQL syntax, keeps track of our users for us. Every user has a unique numeric ID, stored in the id column. (The special SERIAL datatype in PostgreSQL ensures that each row we INSERT into the database will have a unique value for id.) For the purposes of this column, we ignore the security issues associated with storing passwords in plain text.

Next, we define three columns of type TEXT, which PostgreSQL uses to define limitless text fields. Each of these fields is also given an integrity check to ensure that the value can be neither blank nor NULL. We also define the username column to be UNIQUE.

Now, from the standpoint of data integrity, we already have done our jobs. We can be sure that no two users will have the same user name, that each e-mail address (that is, each person) is allowed to have more than one user name in the system, and that the user name, e-mail address and password cannot be blank. Everything else is just icing on the cake, right?

Well, yes—but only if we are willing to give our users database errors. Most of us would prefer to offer a softer landing to our users, telling them not only that (for example) their chosen user name already has been taken by someone else, but also shielding them from the errors that PostgreSQL displays.

This means our application is going to need to take the user's requested user name, check for it in the database, and then either display an error message (prompting the user to try again) or INSERT a new row in the database.

Here is a simple example of how we might do this in a Web page. I use simple CGI programs written in Perl for the server-side examples this month, in no small part because they tend to be simple to understand and try on any host. Listing 1 (register.html) shows the HTML form users will see when they want to register with the site. Listing 2 (register.pl) shows the CGI program that will accept the form contents, check the user name in the database and then produce a message in response. My assumption is that the form will be in the main Web document root directory, and the CGI program will be in the cgi-bin directory. Obviously, if you are using a server-side language, such as PHP, the two can exist side by side.

The HTML form has three text fields, in which the user is requested to enter a user name, a password and an e-mail

Listing 1.
register.html

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">
  <head><title>Register</title>
  </head>
  <body>
    <h2>Register</h2>
    <form action="/cgi-bin/register.pl" method="post">
      <p>Username: <input type="text" name="username" /></p>
      <p>Password: <input type="password" name="password" /></p>
      <p>E-mail address: <input type="text" name="email_address" /></p>
      <p><input type="submit" value="Register" /></p>
    </form>
  </body>
</html>
```

THE PENGUIN

Another server down. Another night at the office. Whether it's deploying a hundred new servers or installing the latest security patch, it doesn't matter. You're sleeping with the servers again. Penguin Computing® introduces Scyld Enterprise *WebMaster*™. Its centrally-managed, highly available architecture makes large pools of Linux servers act like a single, con-

sistent virtual system, significantly reducing complexity and time to deploy servers on the fly. Its highly secure environment makes your server farm a 'virtual fortress' and its simplified manageability cuts your maintenance windows and administration by up to 80%, dramatically improving TCO. So go on home, catch some zzzzs and know that Penguin is standing guard.

MAY CAUSE DROWSINESS



PENGUIN HIGH DENSITY CLUSTER. The highest density, modular blade server architecture on the market. With powerful Scyld Enterprise *WebMaster*™ for single point command and control, and AMD Dual Core Opteron™ for a highly productive user experience.

www.penguincomputing.com



REALLIFELINUX

Penguin Computing and the Penguin Computing logo are registered trademarks of Penguin Computing Inc. Scyld Enterprise *WebMaster*™ and the Highly Scyld logo are trademarks of Scyld Computing Corporation. AMD Opteron and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices, Inc. Linux is a registered trademark of Linus Torvalds. ©2006 Penguin Computing, Inc. All rights reserved.

Listing 2.

register.pl

```
#!/usr/bin/perl

use
strict;
use diagnostics;
use warnings;

use
CGI;
use DBI;
# -----
# # Connect to the database
# -----
my $dbname = 'atf';
my $dbuser = 'reuven';
my $dbpassword = '';

my
$dbh = DBI->connect("DBI:Pg:dbname=$dbname",
                  $dbuser, $dbpassword,
                  {
                    AutoCommit => 1, RaiseError => 1,
                    PrintError => 1, ChopBlanks => 1}) ||
    print "<p>Error connecting: '$DBI::errstr' </p>";

# -----
# CGI startup
# -----
my $query = new CGI;
print $query->header("text/html");
print $query->start_html(-title => "Site registration");

my $username = $query->param("username");
my $password = $query->param("password");
my $email_address = $query->param("email_address");

# -----
# Check the parameters
# -----
my @missing_data = ();

push @missing_data, "The username"
    unless $username;

push @missing_data, "A password"
    unless $password;

push @missing_data, "The e-mail address"
    unless $email_address;

if (@missing_data)
{
    foreach my $missing_field (@missing_data)
    {
        print "<p>Sorry, but you are missing:</p>\n";
        print "<ul>\n";
        print "<li> $missing_field</li>\n";
        print "</ul>\n";
        print "<p>Please back up and try again.</p>\n";
        exit;
    }
}

# -----
# Try to register the user
# -----
my $select_sql = "SELECT COUNT(*) FROM Users WHERE username = ?";
my $select_sth = $dbh->prepare($select_sql);
$select_sth->execute($username);

my ($username_is_taken) = $select_sth->fetchrow_array();

# Is this username taken? If so, give an error
if ($username_is_taken)
{
    print "<p>Sorry, but the username '$username' was already taken.
Please back up and try again.</p>\n";
}

# Otherwise, insert the new trio into the
database
else
{
    my $insert_sql = "INSERT INTO Users (username, password,
                        email_address)
                        VALUES (?, ?, ?)";
    $dbh->do($insert_sql, {}, $username, $password, $email_address);

    print "<p>Added the username '$username' to the system!</p>\n";
}

print $query->end_html;
```

address. (I personally prefer to use e-mail addresses as user names, but I realize many people don't, so I'll add it this month.) We know from the database definition that the user name must be unique in the system, but the name and e-mail address can each exist multiple times.

The registration program (in Listing 2) is a relatively simple CGI program. It connects to the database using Perl's DBI interface, then uses the standard start of CGI programs, grabbing parameters and generally getting ready. The program then checks the database to see if the user name already exists, returning the

number of rows it matches in the database. If no rows match, we can assume that the user name is available. (There is something of a race condition here, but we're not going to complicate things with transactions for this small example.)

Making the Form Dynamic

This is the type of registration form with which most of us are familiar. Moreover, this is the type of registration form many of us continue to implement on various sites. It's easy for the programmer to build, it's easy to understand and debug, and it's



One

PGI Unified Binary™

Now, PGI® compilers can generate a single PGI Unified Binary executable fully optimized for both Intel EM64T and AMD64 processors, delivering all the benefits of a single x64 platform while enabling you to leverage the latest innovations from both Intel and AMD. PGI Fortran, C, and C++ compilers deliver world-class performance and a uniform development environment across Linux and Windows as part of an integrated suite of multi-core capable software development tools. Visit www.pgroup.com to see why the leading independent software vendors in structural analysis, computational chemistry, computational fluid dynamics and automotive crash testing choose PGI compilers and tools to build and optimize their 64-bit applications.



The Portland Group™
www.pgroup.com ++ 01 (503) 682-2806

compatible with all browsers out there.

The problem is not in the technical underpinnings of the program, but rather in the user interface. From nontechnical users' perspectives, it doesn't make sense for them to enter a user name, password and e-mail address, then submit it, and only then find out that the user name is unacceptable. Surely there must be a way to fix this!

The only way for forms to be checked before they are submitted to a server is to use a client-side language—that is, a language embedded within the Web browser, which can attach itself to browser window events. The universal standard for such a language is ECMAScript, because it was ECMA International (formerly known as the European Computer Manufacturers Association) that approved and published the standard. However, most people refer to ECMAScript by the

language that inspired the standard, namely JavaScript.

JavaScript is almost always found within the pages of an HTML document. We can define and invoke functions inside of the document, triggering the invocations with event handlers. Thus, we can check the contents of a form when someone clicks the submit button, before the contents are sent to the server. We can change styles when the mouse moves on (or off) particular text and graphics. And, we can execute functions when someone enters or exits an HTML form element.

Listing 3 contains `js-register.html`, a modified version of `register.html`. The basic idea in this file is that as soon as the username text field is modified by the user, the browser executes the `checkUsername` function.

This is the way that most client-side Web programs are structured. JavaScript functions do the actual work, but they

Listing 3.

js-register.html

```

<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">
  <head><title>Register</title>

  <script type="text/javascript">
<!--
    var usernames = ['abc', 'def'];

    function removeText(node) {
      if (node != null)
      {
        if (node.childNodes)
        {
          for (var i=0 ; i < node.childNodes.length ; i++)
          {
            var oldTextNode = node.childNodes[i];
            if (oldTextNode.nodeValue != null)
            {
              node.removeChild(oldTextNode);
            }
          }
        }
      }
    }

    function appendText(node, text) {
      var newTextNode = document.createTextNode(text);
      node.appendChild(newTextNode);
    }

    function setText(node, text) {
      removeText(node);
      appendText(node, text);
    }

    function checkUsername() {

      var new_username = document.forms[0].username.value;
      var found = false;
      var warning = document.getElementById("warning");

      var submit_button = document.getElementById("submit-button");

      // Is this new username already taken? Iterate over
      // the list of usernames to be sure.
      for (i=0 ; i<usernames.length; i++)
      {
        if (usernames[i] == new_username)
        {
          found = true;
        }
      }

      // If we find the username, issue a warning and stop
      // the user from submitting the form.
      if (found)
      {
        setText(warning, "Warning: username '" + new_username
+"'" + " was taken!");
        submit_button.disabled = true;
      }

      else
      {
        removeText(warning);
        submit_button.disabled = false;
      }
    }
--> </script>
</head>
<body>
  <h2>Register</h2>
  <p id="warning"></p>
  <form action="/cgi-bin/register.pl" method="post">
    <p>Username: <input type="text" name="username"
onchange="checkUsername()" /></p>
    <p>Password: <input type="password" name="password" /></p>
    <p>E-mail address: <input type="text" name="email_address" /></p>
    <p><input type="submit" value="Register" id="submit-button"/></p>
  </form>
</body>
</html>

```

are invoked by event handlers defined in the HTML. So, in Listing 3, we see:

```
<p>Username: <input type="text" name="username"
  onchange="checkUsername()" /></p>
```

This tells the browser that when the username text field changes, it should invoke the checkUsername() function. When this function is executed, it begins with the following:

```
var new_username = document.forms[0].username.value;
```

The new_username variable gets the value of the username text field. We do this by starting off with the document object (representing our HTML document), then taking the first element of its forms array (representing the first, and only, form in the document). The username property of the form gives us the node for the username text field, whose value we can then retrieve (as a string) with the value property.

Traversing the tree in this way is typical when working with JavaScript. However, it is also possible to jump immediately to a particular form element, assuming that element has been assigned an id attribute. IDs must be unique within a document, meaning that we can find a node with the appropriate method:

```
var warning = document.getElementById("warning");
var submit_button = document.getElementById("submit-button");
```

Each of the above two lines uses document.getElementById to retrieve a node from the document tree, identified with an id attribute. (If nothing matches, the variable is set to the null value.)

The list of user names has been hard-coded in Listing 3, which is something you would never try in an actual application. I discuss this further below, and we will find production-quality, Ajax-style solutions next month.

Dynamically Modifying the Page

Now that we have a list of user names in JavaScript, we want to force the user to choose a user name that does not clash with one already in use. We will do this by checking the proposed user name against the list that we have already collected. If the user name is already taken, we will warn the user by modifying the HTML of the current page and then by disabling the submit button. Only when the chosen user name is new and unique will the user be allowed to submit it to the server. This doesn't mean we will remove our uniqueness checks on the server or in the database, but it offloads some of that checking to the client and makes the application more immediately responsive to the user's needs.

We do this by iterating over usernames, the array containing user names:

```
for (i=0 ; i<usernames.length; i++)
{
  if (usernames[i] == new_username)
  {
    found = true;
  }
}
```

If we find a match between the user's requested new user name and one in the array, then we set the found variable to be true. Otherwise, it continues to be set to false. This then

Books on Ajax and JavaScript

In working on these columns, I have found several good books on the subject of HTML, JavaScript, Ajax and related technologies.

The two most comprehensive books on the subject are O'Reilly's *JavaScript: The Definitive Guide* by David Flanagan and *Dynamic HTML: The Definitive Reference* by Danny Goodman. There is a fair amount of overlap between these two books, and they are definitely meant to be reference books rather than tutorials. That said, experienced Web developers interested in learning about client-side programming probably will learn a great deal from these books. And once you're experienced, you will undoubtedly use these two books often, checking everything from the cross-platform compatibility of various JavaScript objects to how JavaScript interfaces with the DOM.

Newer and less experienced Web developers would probably do well to start with a gentler introduction to these technologies. One of the best, and funniest, that I've seen is O'Reilly's *Head Rush Ajax* by Brett McLaughlin. My one criticism is touted as one of the book's strengths—namely, that it presents the same information in many different ways to ensure that you will remember it. The book might be a bit annoying for experienced Web developers who will want to get to the meat more quickly, and who might be frustrated by the repetition. Nevertheless, I think that this is a worthwhile read for anyone starting in the Ajax world.

A middle-of-the-road book that might appeal to more experienced Web developers, while providing a tutorial and introduction to many JavaScript concepts described here, is *Professional JavaScript for Web Developers* by Nicholas Zakas and published by Wrox. I didn't like the way in which Zakas used his own (freely available) JavaScript library throughout the book, but I did think that the examples and explanations were carefully chosen and interesting, and they also helped to illuminate some of the shadier sides of JavaScript. Zakas contributed to another Wrox book, *Professional Ajax* by Zakas, Jeremy McPeak and Joe Fawcett, which I found to be less integrated and less enjoyable than either the JavaScript book or the *Head Rush Ajax* book from O'Reilly.

tells us whether we need to warn the user about a conflict and disable the submit button or vice versa.

Warning the user consists of two steps. The first involves setting the warning text—that is, the text inside of the <p> tags above the form—to an appropriate message. We already set the variable warning to point to that node at the beginning of the checkUsername function, which means that we now must eliminate all children of the warning node. Actually, we don't want to eliminate all children, but merely the ones with a nodeValue property, because that is where text is stored. The removeText function does that by iterating over each of a node's children, checking to see if it contains text and removing it if it does:

```
if (node.childNodes)
{
  for (var i=0 ; i < node.childNodes.length ; i++)
  {
    var oldTextNode = node.childNodes[i];
    if (oldTextNode.nodeValue != null)
    {
      node.removeChild(oldTextNode);
    }
  }
}
```

Once we have removed text children from the warning node, we then can add a new text child to the warning node, containing the message we want to display. This is done in the `appendText` function:

```
function appendText(node, text) {
    var newTextNode = document.createTextNode(text);
    node.appendChild(newTextNode);
}
```

As of this point, the user has received a warning about the chosen user name, indicating that it will not be accepted because the user name was already taken. However, we cannot rely on users to read and follow the instructions in a warning message. Rather, we should disable the form's submit button, making it difficult for users even to send the bad user name to our server-side program. We can do this by setting the submit button's disabled property:

```
submit_button.disabled = true;
```

To recap—when the user enters a value in the username text field that is in the usernames array, we remove any existing text children from the warning node. We then add a new text child node to warning, indicating that the chosen user name already has been taken. Finally, we disable the submit button in the HTML form.

Of course, we want the user to submit the form eventually, but only after entering a user name not in the usernames array. This means we must remove text children from the warning node, and then re-enable the submit button:

```
removeText(warning);
submit_button.disabled = false;
```

Sure enough, this combination of JavaScript functions seems to do the trick. User names that are not in the usernames array remove any error messages and re-activate the form, allowing us to submit it to the server-side CGI program and register with the site. User names that are in the array, however, produce a warning and stop us from being able to submit the form. It's not Ajax just yet, but it is more responsive to the user than our pure server-side solution.

Considerations

Of course, the program in Listing 3 is fatally flawed in several ways. The biggest, by far, is the fact that the usernames array is hard-coded in the JavaScript. It goes without saying that hard-coding a list of user names in this way is guaranteed to fail, because the list of users is stored in a database table, and we have not connected the database with the program.

We could overcome this problem by generating the usernames array from the database. In other words, our server-side program would create part of our client-side JavaScript program dynamically. Thus, instead of what we see in Listing 3:

```
var usernames = ['abc', 'def'];
```

we would use a server-side program to do something like

the following:

```
my $output = "[";
my $sql = "SELECT username FROM Users";
my $sth = $dbh($sql);
$sth->execute();
while (my ($username) = $sth->fetchrow_array())
{
    $output .= "'$username', ";
}
$output .= "];";
```

We would then insert `$output` into the resulting HTML file, ensuring that the value of usernames would have the most complete and up-to-date list of user names in the system.

But even this is likely to cause serious security concerns in a production application, because it means that every user name in your system—including those with poorly chosen passwords—will be available to everyone visiting your registration page, simply by looking at the HTML source code.

Although it is true that every user name has a password, and that someone would have to guess the password associated with a user name in order to break into your system, can you really vouch for the quality of every password? Moreover, the user names themselves might be clues as to the number or types of users on your system. In short, you really don't want a production system to list the user names for a potential attacker, secure as you might believe your system to be.

There is also an efficiency problem here. As your list of users grows, the length of the usernames array will grow as well. Can you imagine the time it would take to generate and download the JavaScript for a site with 10,000 users?

The solution to all of these problems is, of course, Ajax. Rather than checking the proposed new user name against an array in our JavaScript application, we will have JavaScript submit the proposed user name to the server, find out whether it already has been taken and act accordingly—all without forcing the user to switch to a different page of HTML! This is the underlying magic that makes Ajax applications so compelling; they keep you on the same page longer than traditional Web applications, thus providing a smoother user experience.

Conclusion

We're making some progress on our way to Ajax heaven. We now have an application—user registration—for which old-style Web development provides an answer, but one that feels clunky to the user. The solution we saw in this month's column works well, but requires that the JavaScript contain a usernames array with all user names on the system. For performance and security reasons, this is a bad idea, and we should look for a different solution. Next month, we will start to look at a genuine Ajax solution to this problem, making our application look and feel smoother, while increasing its security as well. ■

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 atnoland.lerner.co.il.



Just because your IT equipment goes dark
doesn't mean you have to go blind.



SecureLinX SLC

Reach your IT equipment
from anywhere as easily as
changing a light bulb.

SecureLinX™ SLC secure console
managers from Lantronix
provide consolidated access so
you can control, manage and
repair your IT equipment from
anywhere, at anytime.

For your
free Console
management
white paper visit
lantronix.com/slcwp/

Network down?
No sweat.
On the road?
No problem.

With SecureLinX you can finally
achieve true "lights out," out-of-
band data center management.
And SecureLinX SLC has your
back with the highest level of
security available. For more
information, check out the
specs at lantronix.com/slc/
or call us at (800) 422-7055.



LANTRONIX®
Network anything. Network everything.™

© Lantronix, 2006. Lantronix is a registered trademark,
and SecureLinX is a trademark of Lantronix, Inc.



Lantronix's SecureLinX SLC16
Awarded Network Computing's
Editor's Choice Award





MARCEL GAGNÉ

Your Media: Out of the Kitchen and into Every Other Place

No question. Your taste in music is certainly among the best. Why not share the tracks that make your heart sing with others?

This is certainly a fascinating collection of music, François! I hadn't realized you were such a big fan of the Indie music scene. You've got some great tunes queued up in Amarok, and there's plenty here that I've never even heard of. You should let me know about it, so I can check out some of it myself.

Quoi? No, François, I don't want you to e-mail me your current playlist. There are other, more interesting ways to do this now. Sharing your musical taste with others is becoming quite popular, *mon ami*. People have "now playing" or "recently played" lists in their blogs, and music discussion is a staple of social networking sites all over the Web. Furthermore, you already are running some great Linux software that will let you share your musical tastes with one, two, five or a thousand people, if you so desire. Of course, I will show you, but for now, we must get ready. As you can see, our guests have already arrived.

Welcome, everyone, to *Chez Marcel*, where we match the best in Linux and open-source software with superb wines and, of course, the best of clientele. Please, find your tables and make yourselves comfortable while my

faithful waiter attends to the *vin*. François, please head down to the wine cellar. I think the 2004 Village Latour Côte du Rossillon from Maison Chapoutier would make an excellent counterpoint to today's menu. Why not marry excellent musical taste with an excellent and flavorful red wine? You'll find a few cases in the third row of the south wing. *Vite*, François! Our guests must be parched.

François and I were discussing the joys of sharing music we enjoy with friends (and anyone else who shares an interest for that matter). Your Linux system makes this easy with the right programs. One way to share your entertainment, *mes amis*, is simply to tell the world. Kopete, the multiprotocol instant messaging client distributed with KDE, has a neat little feature that lets you do just that. Say you happen to be logged in to an IRC channel or having a chat with a friend on-line, using Jabber, Yahoo or some other service. Typing `/media` in the chat field at the bottom of your chat window is all it takes. The system then transmits a message telling your chat partners the artist, title and album of whatever song you are listening to (Figure 1).

Ah, François, you have returned. Excellent. Please, pour for our guests. I envy you, *mes amis*; you will truly enjoy this wine.

Sending your current track information via Kopete works with Amarok, Kaffeine, Noatun and other KDE applications. It also works well with XMMS. In order to get Kopete to transmit this information, however, you may need to turn on the plugin that does the job. To do this, click on Settings in Kopete's menu bar and select Configure Plugins. A window pops up with a list of Kopete's available plugins in the left-hand sidebar (Figure 2).

Make sure the Now Listening check box is checked. Now, look over on the right, and you'll see three tabs. You can change the format of the display to your instant messaging client if you don't happen to like the current setting. The interesting tab here is the one labeled Advertising Mode. The default requires that you enter `/media` to transmit the information, but you can make it automatic here. Another setting transmits the Now Listening information in your on-line status information. When you are happy with the settings, click OK.

Of course, this is only text information, and no one but yourself can hear what you are listening to. To get the sound out to others, you need to do a little broadcasting. A common scenario for a lot of people at home is having a huge collection of music on one PC but two or more computers in other parts of the house. Wouldn't it be great to be able to play music on one PC and have it broadcasted on the other computers,

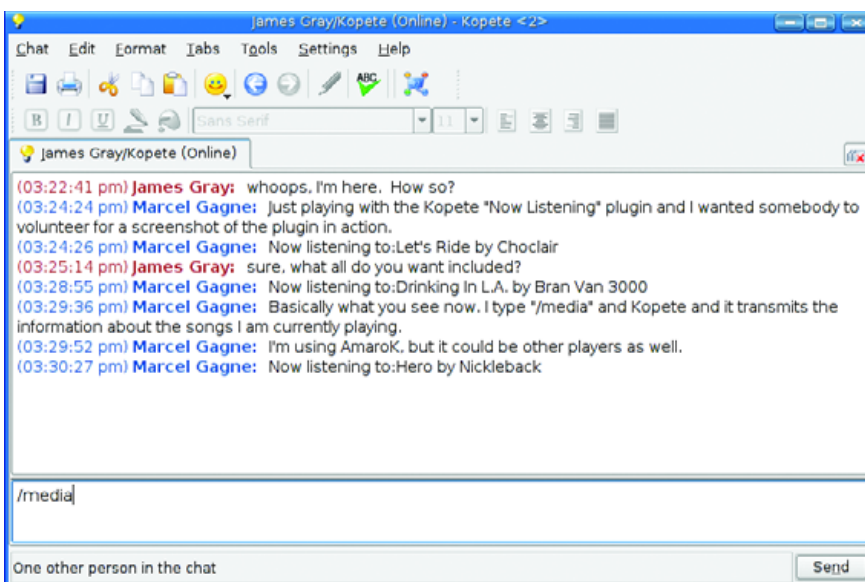



Figure 1. With Kopete's Now Listening plugin, your favorite music player can tell your friends what you are listening to.

Real performance, virtual flexibility.

Get the Dual-Core Intel® Xeon® Processor in DCX9000 series system and enjoy great performance for your 32-bit and 64-bit HPC applications plus the flexibility of Intel® Virtualization Technology. With Intel's advanced virtualization support, you increase availability and reduce complexity.



ZT Dual Core Server DCX9017

2 x Dual-Core Intel® Xeon® Processor 5050

(2x 2MB L2 Cache, 3 GHz, 667MHz FSB, Intel® HT, EM64T)

- Intel® 5000V Chipset Server Board
- 1GB DDRII 533 Fully Buffered DIMM (Up to 16GB)
- 2 x Seagate® 320GB SATAII 16MB Cache Hard Drive (RAID 1)
- 4 x 1" SAS / SATA Hot-swap Drive Bays (with SES2 Support)
- 16x Dual-layer DVD-RW and Floppy Drive
- 4 x SATAII Ports via ESB2 Controller (RAID 0, 1, 5, 10 support)
- Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
- Mid Tower Chassis w/ 645W Low Noise Power Supply
- SuperDoctor III Server Management Software
- 3-Year Limited Warranty + 1-Year On-site Service

ZT Dual Core 2U Server DCX9018

2 x Dual-Core Intel® Xeon® Processor 5060

(2x 2MB L2 Cache, 3.20 GHz 1066MHz FSB, Intel® HT, EM64T)

- Intel® 5000V Chipset Server Board
- 1GB DDRII 533 Fully Buffered DIMM (Up to 16GB)
- 2 x Seagate® 320GB SATAII 16MB Cache Hard Drive (RAID 1)
- 6 x 1" Hot-swap SAS / SATA Drive Bays
- Slim DVD-ROM and Floppy Drive
- 4 x SATAII Ports via ESB2 Controller (RAID 0, 1, 5, 10 support)
- Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
- 2U Rackmount Chassis w/ 550Watt Power Supply
- SuperDoctor III Server Management Software
- 3-Year Limited Warranty + 1-Year On-site Service

ZT Dual Core 3U Server DCX9019

2 x Dual-Core Intel® Xeon® Processor 5130

(4MB L2 Cache, 2 GHz, 1333MHz FSB, Intel® EM64T)

- Intel® 5000V Chipset Server Board
- 2GB DDRII 533 Fully Buffered DIMM (Up to 16GB)
- 2 x Seagate® 73GB 15,000rpm SAS Hard Drive
- 16 x 1" Hot-swappable SAS / SATA Drive Bays
- Slim DVD-ROM and Floppy Drive
- 8 Channel SAS RAID Controller (RAID 0, 1, 10 support)
- Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
- 3U Rackmount Chassis Server Chassis w/ 800W Redundant Power Supply
- SuperDoctor III Server Management Software
- 3-Year Limited Warranty + 1-Year On-site Service

\$1,799

\$2,299

\$3,999

- Highest quality server for smooth-running applications
- Enterprise storage solutions for system reliability and stability
- Flexibility from SATA to SAS drives

Go to

ztgroup.com/go/linuxjournal

Call

866-ZTGROUP (866-984-7687)

Promote code : Lj1006



Purchaser is responsible for all freight costs on all returns of merchandise. Full credit will not be given for incomplete or damaged returns. Absolutely no refunds for merchandise returned after 30 days. All prices and configurations are subject to change without notice and obligation. Opened software is non-refundable. All returns have to be accompanied with an RMA number and must be in re-sellable condition including all original packaging. System's picture may include some equipments and/or accessories, which are not standard features. Not responsible for errors in typography and/or photography. All rights reserved. All brands and product names, trademarks or registered trademarks are property of their respective companies. Celeron, Celeron Inside, Centrino, Centrino Logo, Core Inside, Intel, Intel Logo, Intel Core, Intel Inside, Intel Inside Logo, Intel SpeedStep, Intel Viiv, Itanium, Itanium Inside, Pentium, Pentium Inside, Xeon and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

As you play more and more songs, the site starts to create a profile of your musical tastes, then adds neighbours, stations of people whose musical tastes mirror your own.

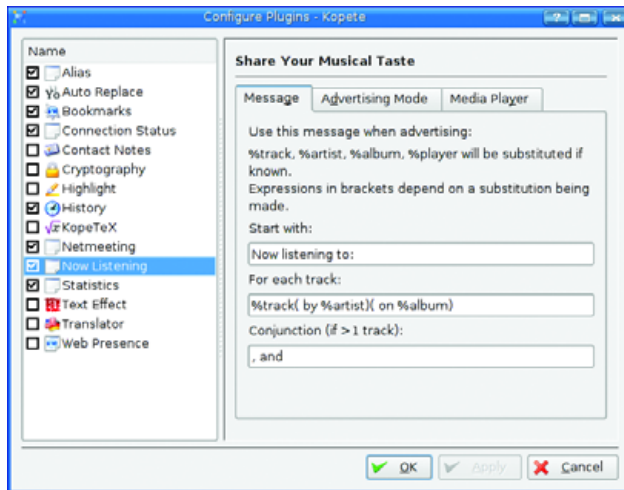


Figure 2. Activating and Configuring Kopete's Now Listening Plugin

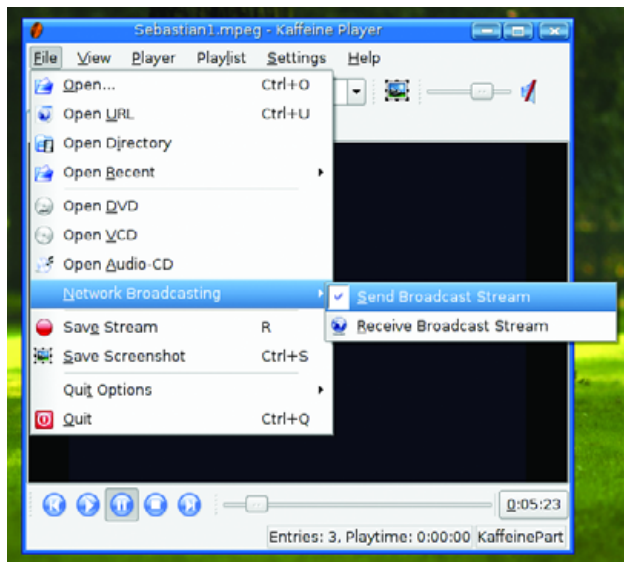


Figure 3. It's easy to turn Kaffeeine into a broadcasting media player.

wherever they might be in the house? Becoming a local broadcaster on your network is easier than you might think. In fact, you may have a few programs that do the job nicely already installed on your system.

One of these is Jürgen Kofler and Christophe Thommeret's Kaffeeine, a popular KDE media player. Kaffeeine is generally thought of as a video player, but it also can serve up DVDs, VCDs, audio CDs or a variety of multimedia files stored on your system. Some Linux distributions that include the Kaffeeine player also embed it in Konqueror for viewing on-line videos in Web pages. Kaffeeine also has a broadcast feature. Click File on Kaffeeine's menu bar, then navigate to Network Broadcasting, and select Send Broadcast Stream (Figure 3).

A small window appears asking you to select a port number. By default, the port already is selected as 8080, but you can override it here. Click OK to close the dialog, and you are ready to broadcast. In fact, anything you play,




Figure 4. On the receiving end, you must tell Kaffeeine which host is doing the broadcasting.

whether it is a music track or a video, is now being broadcast. Other people on your network who want to pick up this broadcast need to fire up their own copy of Kaffeeine and do a slight change to their configuration. Click File on the menu bar and head to the Network Broadcasting submenu, but this time, choose Receive Broadcast Stream. A configuration dialog appears (Figure 4).

Enter the sender's IP address, the port on which the sender is transmitting (the one you set up in the Kaffeeine program doing the broadcasting), and click OK. Just like that, you are watching (or listening) to whatever the local DJ or VJ in your home is sending out.

Several "organic" music services have appeared on-line recently as part of the whole Web 2.0 mania. Some of them, I confess, are quite fascinating. My favorites to date are Pandora and Last.fm. Pandora is an interesting service that suggests, and plays, musical selections based on songs you tell it you enjoy. It turns out that it is very good at guessing what other music you will like. Pandora is great, but it's a solitary experience. What sets Last.fm apart is the social networking aspect. You can discuss music with friends, blog about your favorite tracks, discover new music and find out what your friends are listening to. And, of course, you can tell them what you are listening to on your Linux system. To get in on the fun, you first need a Last.fm account. Then, using a cool little piece of software called Audioscrobbler, your favorite Linux music player can tell Last.fm what you are listening to. Audioscrobbler, by the way, is essentially a database system that tracks listening habits and generates statistics used to predict likes, dislikes and so on.

In the December 2005 issue, I told you about the incredibly amazing Amarok, the KDE jukebox music player. This program is a must-have for music lovers and remains my favorite media player. Amarok's features are too numerous to list here, but let me give you a recap. There's a powerful cover manager (downloads covers from Amazon), a context browser that keeps track of your favorite and most-listened-to songs, a skinnable interface, iPod support (other players work as well), great visualization tools (using libvisual) and more. There's even a lyric

NOTE:

You can check out Marcel's article on Amarok by visiting the *Linux Journal* Web site at www.linuxjournal.com/article/8558.

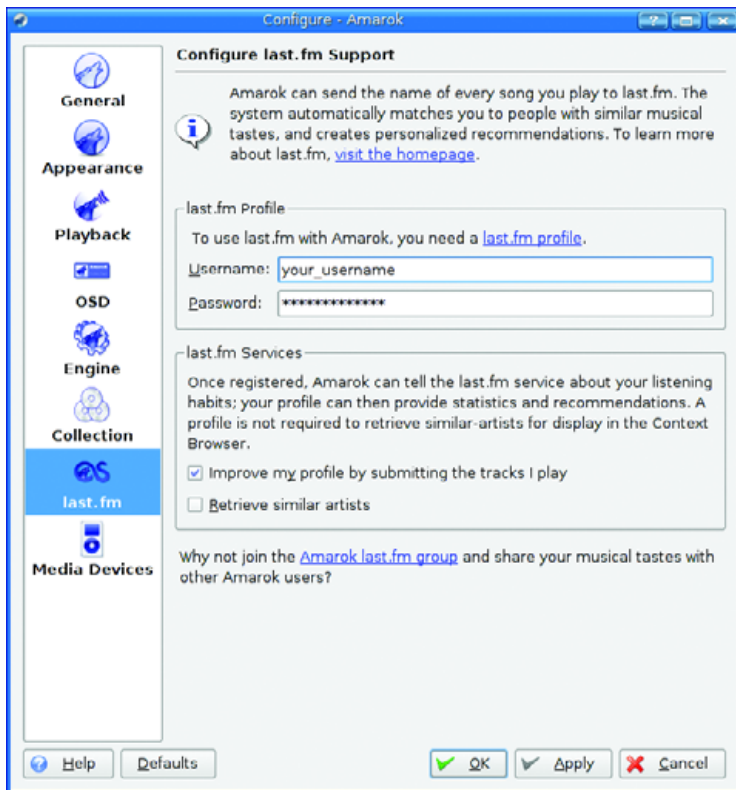


Figure 5. The latest Amarak allows you to broadcast your musical tastes via Last.fm.

download feature so you can sing along with your favorite tunes without worrying about whether you are getting the words right. Now, the latest version of the amazing Amarak also features Last.fm support with built-in Audioscrobbler support.

To use Amarak with Last.fm, click Settings on the menu bar and select Configure Amarak. The configuration window appears (Figure 5) with various categories listed in the left-hand sidebar. One of these is for the Last.fm service, courtesy of the Audioscrobbler program. Enter your Last.fm user name and password, then click OK.

On the GNOME side of things, there also has been some recognition of Last.fm's popularity with music lovers. That's reflected in the excellent Banshee music player. To start telling the world what you are listening to with Banshee, click Edit on the menu bar, and select Plugins. The Banshee plugins configuration window appears with the Audioscrobbler plugin selected. Click the Configuration tab (Figure 6). Click the check box that reads Enable song reporting, enter your user name and password, then click Close.

Whether you choose to use Amarak or Banshee (or something else), the information regarding your musical tracks will now be transmitted to the Last.fm community and people can find out what you are listening to. Earlier on, I mentioned that people could listen to what you are playing as well, and that also is something you can do at Last.fm. Alas, to become a Last.fm DJ, you need to subscribe to the service. To take this service out for a spin, I spent my \$3 US and added a personal radio station to my wftl account. With the paid account, your friends can tune in to your radio station using the free, GPLed, Linux music player distributed on the site (yes, versions are available for other operating systems).

To use the player, unpack the bundle into the directory of your choice (tar -xjvf LastfmLinux-1.1.4.tar.bz2). There is nothing to compile



ASA COMPUTERS

Want your business to be more productive?
 The ASA Servers powered by the Dual Core Intel® Xeon™ Processor provides the quality and dependability to keep up with your growing business.

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

Dempsey/Woodcrest Server \$2024

- 1U Dual Core 5030 CPUs.
- 4GB FBDIMM Memory.
- Supports upto 64GB FBDIMM.
- 120 GB hotswap hard drive.
- 2xintegrated Dual 10/1000 LAN.



Dual Xeon Server starts at \$4015

- 5U Dual Xeon 2.8 Ghz 800 FSB.
- iSCSI or NAS Software options.
- 8x120 GB SATA Hard drives -Upto 18TB.
- 512 MB RAM
- Fail hard drive LED indicator.



Dual Xeon 800FSB Storage starts at \$6,644

- 8U Dual Xeon 2.8 Ghz 800FSB CPUs.
- 2TB of storage (36TB max).
- 1GB RAM
- NAS or iSCSI software options
- 2 x 10/100/1000 Gigabit LAN.



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



"Your Logo Here"

ASA Collocation
 \$75 per month for 1U Rack - 325 GB/month

Storage Solutions
 IDE, SCSI, Fiber Raid solutions
 NAS, DAS, iSCSI, SATA, SAS
 3Ware, Promise, Adaptec
 JMR, Kingston/Storcase solutions

ASA Collocation Special
 First month of collocation free.*

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

All systems installed and tested with user's choice of Linux distribution (free). ASA Collocation—\$75 per month



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



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

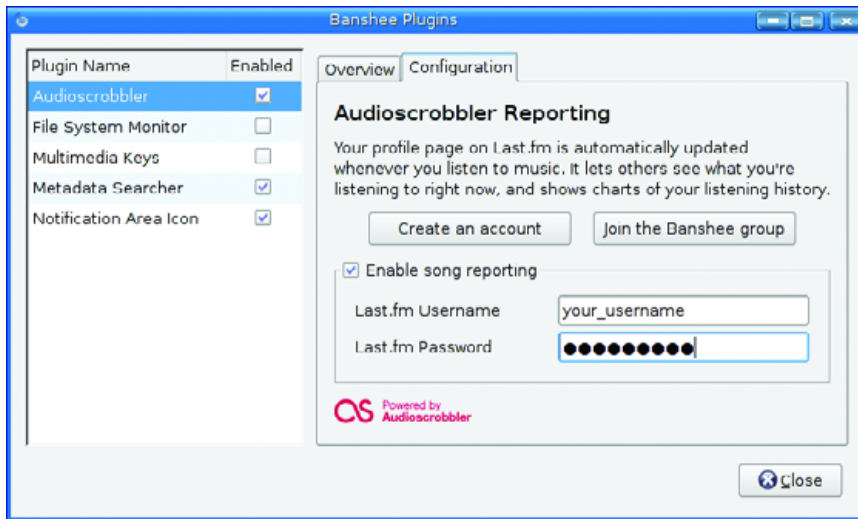


Figure 6. Banshee Configuration Tab

here, so running it is as easy as switching to the directory into which you have extracted it and running the player (`./player`). When you run the player the first time, the setting dialog appears. This is where you enter your Last.fm account information.

Once all the information is entered, you are ready to go. For your friends or family to tune in to the songs you are playing, they need to fire up their player by switching to the directory (or folder) and running the program. In addition to the player, a window labeled Radio Control appears. Before you can play a station, other than your own, you need to create, or add, a radio station to the list on your personal Last.fm page. Once done, you can click the Personal Radio button, and the player takes over, streaming whatever content is coming from your media player, whether it is Amarok, Banshee or anything else that uses Audioscrobbler.

If you would prefer to use a player other than this one, you can do that by clicking the settings button (the tool button on the lower right) and making your choice there. You also can specify a station by entering its address (for example, `lastfm://user/wftl/personal`).

Time appears to be running out, *mes amis*, but I do want to direct your attention momentarily back to Amarok. Once you are part of the Last.fm social network, you don't actually need to use the Last.fm player. Amarok can tune in directly to what Last.fm calls Neighbour Radio. As you play more and more songs, the site starts to create a profile of your musical tastes, then adds neighbours, stations of people whose musical tastes mirror your own.

On Amarok's menu bar, click Engage, then Play Last.fm Stream, then select Neighbour Radio. Even if you don't have any musical neighbours, you also can choose to listen to Global Tag Radio, which plays music based on the Last.fm community's collectively tagged music. Each of these stations is arranged according to genre, whether it be Rock, Pop, Dance, Rap or whatever. It's a great way to discover new music.

The clock on the wall, *mes amis*, she is telling us that closing time is upon us yet again. The collection of music being

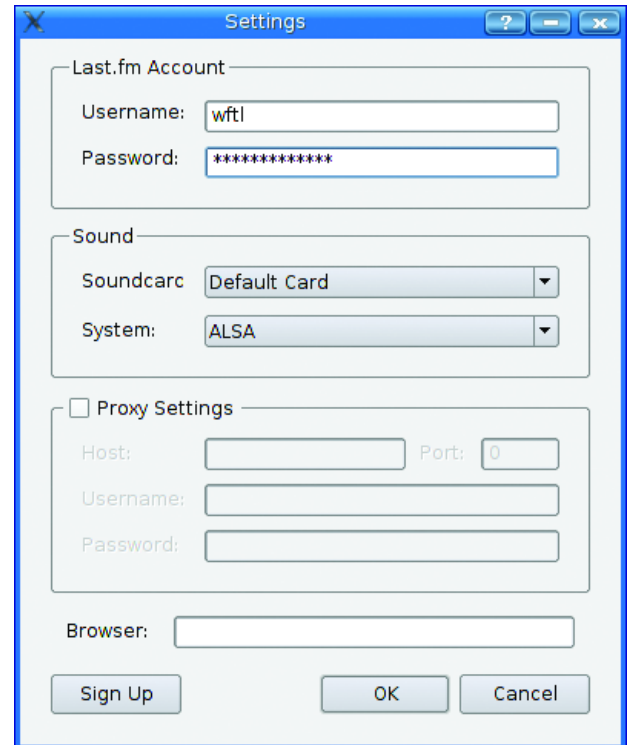


Figure 7. Configuring the Last.fm Media Player



Figure 8. Last.fm's Freely Distributed Music Player

played here in the restaurant at your tables would make for a fascinating stream, indeed. I hear some excellent music—not surprising with a clientele of such impeccable taste. François, please refill our guests' glasses a final time so that we may raise a toast. Please raise your glasses, *mes amis*, and let us all drink to one another's health. *A votre santé! Bon appétit!* ■

Resources for this article: www.linuxjournal.com/article/9172.

Marcel Gagné is an award-winning writer living in Mississauga, Ontario. He is the author of the all new *Moving to Ubuntu Linux*, his fifth 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 maggagne@salmar.com. You can discover lots of other things (including great Wine links) from his Web site at www.marcelgagne.com.

TIP:

Audioscrobbler plugins are available in the downloads section of Last.fm for XMMS, Noatun and other Linux players.

SUPERMICRO®

2U in 1U 4+1 Add-on Cards

SuperServer 6015X-8/T

- Dual Intel® Dual-Core 64-bit Xeon™ processors support
- 1333/1066/667MHz Front-Side Bus
- 2 PCI-E x8 w/RSC-R1UEP-2E8 (full-length)
1 Universal PCI-X 133MHz / PCI-E x8 (full-length)
1 Universal PCI-X 100MHz / PCI-E x4 (low-profile)
optional 2 PCI-X 133MHz w/RSC-R1UEP-A2X (full-length)
- 1 SIM1U management slot (optional 3rd data LAN)
- Up to 32GB Fully Buffered DIMM Memory 667/533 MHz
- 700W high-efficiency power supply w/I²C management
- 5x heavy duty counter-rotating fans w/optimal fan speed control and air shroud
- 3 hot-swap SCA/SATA drive bays
- Dual Gigabit LAN ports



SUPER Serverboard X7DBX-8/T

Supermicro's Feature Advantages

- 4 Universal Expansion cards support PCI-X and PCI-E with interchangeable configurations
- Flexible Supermicro Intelligent Management (SIM): IPMI 2.0 with Built-in Virtual Media over LAN, optional KVM over LAN and 3rd LAN port
- High Efficiency Power Supply (85% and up) reducing energy costs and customer TCO
- Optimized air-flow for efficient cooling
- SAS version available



Dual-Core Xeon 5100/5000 Series

AMAX Corp.
1-800-800-6328
www.amax.com

Arrow Electronics
1-888-427-2250
www.arrowncp.com

ASI
1-800-2000-ASI
www.asipartner.com

Bell Micro.
1-800-232-9920
www.bellmicro.com

Ingram Micro
1-800-456-8000
www.ingrammicro.com

MA LABS
1-408-941-0808
www.malabs.com

Synnex Inc.
1-800-756-5974
www.synnex.com

Tech Data
1-800-237-8931
www.techdata.com



DAVE TAYLOR

Analyzing Log Files

Ever wondered what your Web server is doing, but find that you don't have a stats or analytics package installed? In fact, analyzing log files is a perfect task for the Linux command line and, by extension, shell scripts too.

If you're running Apache, and you probably are, you've got a file called `access_log` on your server, probably in `/etc/httpd` or some similar directory. Find it (you can use `locate` or `find` if needed).

First, let's see how many hits you've received—that is, how many individual files have been served up. Use the `wc` program to do this:

```
$ wc -l access_log
83764 access_log
```

Interesting, but is that for an hour or a month? The way to find out is to look at the first and last lines of the `access_log` itself, easily done with `head` and `tail`:

```
$ head -1 access_log
140.192.64.26 - - [11/Jul/2006:16:00:59 -0600]
  ↳ "GET /favicon.ico HTTP/1.1" 404 36717 "-" "-"
$ tail -1 access_log
72.82.44.66 - - [11/Jul/2006:22:15:14 -0600]
  ↳ "GET /individual-entry-javascript.js HTTP/1.1"
  ↳ 200 2374 "http://www.askdavetaylor.com/"
  ↳ sync_motorola_razr_v3c_with_windows_xp_via_bluetooth.html"
  ↳ "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; SV1;
  ↳ .NET CLR 1.1.4322; .NET CLR 2.0.50727)"
```

When you have a long list of data like this, you can figure out the most popular individual occurrences by sorting everything, then using the `uniq` command to figure out how often each line occurs.

These log file lines can be darn confusing, so don't panic if you look at that and become completely baffled. The good news is it's not important to know what every field details. In fact, all we care about is the date and time in square brackets, and the name of the individual file requested after the "GET" line.

Here you can see that the first line in the access log is from 11 July at 16:00:59 and the last line is from 11 July at 22:15:14. Calculate this out, and we're talking a window of about six hours and 15 minutes, or 375 minutes. Divide the total number of hits by this time passage, and we're seeing 223 hits per minute, or a pretty impressive traffic level of 3.7 hits per second.

The Most Popular Files Sent

The second common query is to ascertain which files are requested most often, and that's something we can ascer-

tain with a quick call to `awk` to split that field from the log file lines, then a combination of `sort` and `uniq` with its ever-useful `-c` option.

Let's take this one step at a time.

If you go back to the log file line shown above, you'll find that it's the seventh field that contains that value, meaning we can extract it like this:

```
$ head access_log | awk '{print $7}'
/favicon.ico
/0-blog-pics/itunes-pc-advanced-importing-prefs.png
/0-blog-pics/itunes-pc-importing-song.png
/styles-site.css
/individual-entry-javascript.js
/motorola_razr_v3c_and_mac_os_x_transfer_pictures_and_wallpaper.html
/Graphics/header-paper2.jpg
/Graphics/pinstripebg.gif
/0-blog-pics/bluetooth-razr-configured.png
/0-blog-pics/itunes-pc-library-sting.png
```

When you have a long list of data like this, you can figure out the most popular individual occurrences by sorting everything, then using the `uniq` command to figure out how often each line occurs. Then use `sort` again, this time to sort the data from that, prefaced with the largest numeric value to the smallest.

Here's an intermediate result to help you see what's happening:

```
$ awk '{print $7}' access_log | sort | uniq -c | head
535 /
26 //favicon.ico
6 //signup.cgi
1 /0-blog-pics/MVP-Combo_picture.jpg
2 /0-blog-pics/address-book-import.jpg
4 /0-blog-pics/adwords-psp-bids.png
28 /0-blog-pics/aim-congrats-account.png
28 /0-blog-pics/aim-create-screen-name.png
38 /0-blog-pics/aim-delete-screenname-mac.png
29 /0-blog-pics/aim-forget-password.png
```

All that's left is to sort it by most popular and axe all but the top few matches:


```
$ awk '{print $7}' access_log | sort | uniq -c | sort -rn | head
6176 /favicon.ico
5807 /styles-site.css
5733 /Graphics/header-paper2.jpg
5655 /Graphics/pinstripebg.gif
5512 /individual-entry-javascript.js
5458 /Graphics/marker-tray.gif
5366 /Graphics/help-button.jpg
5363 /Graphics/digman.gif
5359 /Graphics/delicious.gif
5323 /0-blog-pics/starbucks-hot-coffee.jpg
```

The first thing you'll notice is that this isn't pages but graphics. That's not a surprise, because just like most Web sites, my own AskDaveTaylor.com has graphics shared across all pages, making the graphics more frequently requested than any given HTML page.

Fortunately, we can force the results to be HTML pages by simply using the grep program to filter the final results of the filter sequence:

```
$ awk '{print $7}' access_log | sort | uniq -c | sort -rn
| grep "\.html" | head
446 /motorola_razr_v3c_and_mac_os_x_transfer_pictures_and_wallpaper.html
355 /how_to_create_new_screen_names_on_aol_america_online.html
```

```
346 /how_do_i_cancel_my_america_online_aol_account.html
293 /pc_to_sony_psp_how_do_i_download_music.html
206 /how_do_i_get_photos_and_music_onto_my_sony_psp.html
198 /how_do_i_get_my_wireless_wep_password_for_my_sony_psp.html
195 /cant_get_standalone_music_player_to_work_on_myspace.html
172 /convert_wma_from_windows_media_player_into_mp3_files.html
166 /sync_motorola_razr_v3c_with_windows_xp_via_bluetooth.html
123 /how_do_i_create_a_new_screen_name_in_aol_america_online_90.html
```

(Yes, yes, I know that the URLs on this site are ridiculously long!)

Now, finally, I can see that the articles about the Motorola RAR phone, AOL screen names and Sony PSP are the most popular articles on the site. Remember, this is a slice for only about six hours too, so the RAZR article is actually being requested an impressive once a minute or so. Popular indeed!

I'm going to stop here now that you've had a taste of how basic Linux commands can be combined to extract useful and interesting data from an Apache log file. Next month, we'll look at one more statistic: how much aggregate data we've transferred. Then, we'll start looking at how to build a shell script that does these sorts of calculations with ease. ■

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.

Hurricane Electric Internet Services...**Speed, Reliability,** **and 24/7 Support...You Can Finally Vacation Worry Free!**

Flat Rate Gigabit Ethernet

1,000 Mbps of IP

\$13,000/month*

Full 100 Mbps Port

Full Duplex

\$2,000/month


Colocation Full Cabinet

Holds up to 42 1U
servers

\$400/month

Order Today!

email sales@he.net or call 510.580.4190

 he.net

* Available at PAIX in Palo Alto, CA; Equinix in Ashburn, VA; Equinix in Chicago, IL; Equinix in Dallas, TX; Equinix in Los Angeles, CA; Equinix in San Jose, CA; Telehouse in New York, NY; Telehouse in Los Angeles, CA; Telehouse in London, UK; NIKHEF in Amsterdam, NL; Hurricane in Fremont, CA; Hurricane II in Fremont, CA and Hurricane in San Jose, CA



JON "MADDOG" HALL

What's in a Name?

Hackers, who are not always crackers, deserve a little Agape.

I was sanding down the hull of the *Agape*, my sloop that I sail from time to time, trying to get the boat ready for its annual painting when one of the younger Pollywogs, Yury, came up to me and complained that someone had called him a "hacker".

"Why does that bother you?" I asked. "I do not break into other people's systems", Yury said. "I program, I do systems administration and I help people a lot with their computer problems."

I stopped sanding a bit and realized that Yury had confused the term Hacker with Cracker. It is easy to do, as the mass media has done much to mix the two terms.

"No Yury", I said. "A hacker is a fine name to be called, it is someone who really knows programming, has studied it, enjoys doing it and shows exceptional skill at solving difficult problems with elegant code. Usually that code is relatively short, what we call a quick hack."

Seeing me talking to Yury, more Pollywogs started coming up. "Then what is a cracker?" asked one of them. "A cracker", I explained, "is someone that breaks into other people's systems, sometimes to do malicious damage, like destroy their Web site, but often just to show people that their sites or systems are insecure."

I went on to explain that there were two main types of crackers, those who were extremely skilled and often very hard to detect without constant diligence and those who were less skilled, often relying on techniques developed by others, but applied to thousands or tens of thousands of systems over the Internet. Often these less skilled people are nicknamed script kiddies.

I also explained that in a lot of countries breaking into another person's system is against the law, even if it is relatively easy and no damage is done, and if the person is caught, there can be dire circumstances.

"But why do these people break in?" asked Yury, "Sometimes it is for financial gain", I explained, and I continued:

Sometimes it is to steal records or change data that will bring some type of financial benefit to the cracker. Sometimes it is to see what other people are doing or to change their software so people will inadvertently send the crackers information that will be useful to them over time. Sometimes it is just for the challenge of breaking in, for peer respect of having "done the deed".

A friend of mine named Marcelo Marques in Sao Paulo, Brazil, had been talking to the Brazilian FBI. They had empirical information that 80% of computer crimes were being either initiated or helped by teenagers. The FBI spokesman expressed frustration to my friend, because he could not understand why these teenagers were doing these things. Marcelo talked with me about crackers and hackers, and after planning for a while, he and his partners from 4Linux, a consulting firm in Sao Paulo, started a program called Hackerteen (www.hackerteen.com.br), under the premise that if you gave training to young people in how to be a really good systems administrator, programmer and community member, you could turn potential crackers into hackers.

So far, the program has been in existence for two years and has graduated its first class of "Black Belt" students.

"Can anyone become a hacker?" asked Yury. I answered:

Well, there were a lot of young people who dropped out of Hackerteen along the way. The Hackerteen program is not an easy one. It is made up of on-line work, study and even on-site visits where the students are given courses in ethics, entrepreneurship, computer security, networks, programming and collaboration. My personal experience tells me that not every person can do programming. I have taught for too many years to say that every person has the gifts to become a top hacker. Nevertheless, a course set up like this can have a very positive effect on young programmers, and some of these graduates now have jobs with companies. I met with the first class of graduates, and I could see what the class meant to both them and their parents.

"Do you have to go to school or take a course to be a hacker?" asked Jay.

"No", I answered, and I continued:

Lots of hackers I know are self-taught. One of the things that I like about free and open-source software is that there is nothing hidden from those who wish to learn. You can see other people's code and learn from it. There are plenty of books on the library shelves, magazines to read, Internet mailing lists and news groups and even electronic books to learn from. Things are very different now than from when I learned to code.

But some people find it difficult to learn by reading books, and some really like to have a guide or mentor along the way. Some do best with classroom instruction, and some are held back by it. My experience has been that a more structured approach gives a more balanced person, but this does not mean that you have to take a course or go to a school, you just have to define a path for yourself and stick with it. Some people find that difficult.

Ethics are particularly difficult, for what may be ethical for one person may not be ethical for many others. Here is where a lot of people may make the wrong decisions without proper guidance. If people think that ethics in computers are easy, they should attend a session of the Ethics Working Group of the SAGE Executive Committee.

"What about girls?" asked another young Pollywog, sort of crinkling his nose. "Can girls become hackers?" "Definitely", I said, and added:

One of my favorite people of all time, one of the earliest programmers and one of the pioneers of the Cobol language was Rear Admiral Dr Grace Murray Hopper. Dr

Hopper encouraged early programmers to share common segments of code, reducing duplication of effort and errors.

In Malaysia, more than 70% of the people in data processing are women, and not just in “menial” tasks, but with jobs of influence and power. At my job in Bell Laboratories, my supervisor was a very talented woman by the name of Bea Fink. My observations have been that there is no inherent differences between a man and a woman doing a job in the computer field. As a former board member of the USENIX Association (www.usenix.org), I was very proud of the work USENIX did with organizations encouraging women to enter the field and acknowledging the women who contributed to it. I also met many fine women who would indeed fit the definition of hacker through USENIX and organizations like the ACM and IEEE.

Some of the Linux community has recognized the need to encourage women to enter the computer science field and started groups like LinuxChix (www.LinuxChix.org); GNUrias (www.gnurias.org.br), a Brazilian/Portuguese organization; Debian Women (www.women.debian.org);

and ChicasLinux (www.chicaslinux.org), a Spanish organization.

“But how do you know you are a hacker?” asked Yuri. I had to sand the bottom of the *Agape* for a while before answering him. I had to think of all the many fine programmers that I had known over the years, all the people who were master craftspeople. I thought about the lists of hackers that I had seen, and the names on those lists. Yet, I did not remember any of these people referring to themselves, unless it was in a passing comment, as a hacker. Then I knew the answer for my young friend.

“Yury”, I said, “you are probably a hacker when some programmer that you respect as a good programmer calls you a hacker.”


The *Agape* is almost ready to sail. ■

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.

Expert Included.

Our experts draw on years of experience to provide powerful, manageable products. Dave likes the Rackform iServ R242, based on the Dual-Core Intel® Xeon® 5100 Series Processor, because it offers unprecedented performance and advanced reliability features. He knows that IT organizations can use Intel® Virtualization Technology to consolidate their server infrastructure and reduce total costs. With the combination of energy-efficient performance, advanced manageability offered by IPMI 2.0, and increased I/O performance provided by Intel® I/O Acceleration Technology, he thinks you’ll like it too.

When you partner with Silicon Mechanics you get more than groundbreaking Intel performance — you get an expert like Dave.



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.



DOC SEARLS

Breaking the Matrix

You need to free your mind, Neo.

I started working full-time for *Linux Journal* in 1999, when we were headquartered in the Ballard District of Seattle—a place I still miss visiting. One day, we took a break and headed over to a local theater where *The Matrix* was showing. It was early April, and the movie had just come out.

The Matrix immediately became my favorite movie of all time. Which is saying a lot, as I'm now 59 years old, and I've seen a lot of movies.

The Matrix doesn't have the best acting or direction. It did break new ground in special effects, but so did many lesser movies. In its many fight and chase scenes, *The Matrix* and its sequels are derivative of countless other products of Hollywood and Asia. (Yes, they were fun and over the top, if you're into that stuff, but they were still derivative, and some of the scenes go on forever.) What made *The Matrix* powerful and important was its exceptionally strong metaphor for closed systems that become so familiar and comfortable we fail to realize they treat us like cattle—or worse, as disposable batteries.

The Matrix and its two sequels tell a familiar sci-fi story: a small rebel force overthrows an oppressive and overpowering empire. In this case, the setting is a future Earth in which machines and their programs have defeated and enslaved almost the entire human race. In this future world, nearly every human is trapped in a vast power plant where each individual

Far more powerful is a belief, held by nearly everybody in the developed world, that the best markets are captive ones.

serves as a source of electrical power for the machines. The movie's setting—the apparently real world of 1999—is in fact a collective illusion programmed and maintained by machines of vast intelligence.

Meanwhile, a population of free humans lives deep underground in a city called Zion and sends hovercrafts to sneak through abandoned sewers, where they can hack into The Matrix and free human prisoners who correctly suspect that something is not right with the only world they've known. It's not worth going further into the details of the movie. I'd be surprised if more than a handful of *Linux Journal* readers haven't seen it. What matters is what we gain from it.

I gained a realization that The Matrix was a metaphor for marketing. For me this was personal.

When I began writing for *Linux Journal* in 1996 (as a contributing editor), I was fairly new to Linux and to the free software and open-source concepts that Linux embodies. But I had been working in marketing—mostly advertising and PR—for two decades. For a stretch of the late 1980s and early 1990s, Hodskins Simone & Searls was one of the top high-tech advertising agencies in Silicon Valley. When I left the agency, I went on to become a successful marketing consultant. So I knew how the sausage was made.

Oddly, marketing (including advertising and PR) is not as powerful as you might think. Given the extraordinary

inefficiencies involved, the actual influence exerted by marketing (and by advertising and PR in particular), is remarkably small. Even the accountabilities introduced with pay-per-click advertising still involve ratios of "impressions" to clicks that run in the lottery range.

Far more powerful is a belief, held by nearly everybody in the developed world, that the best markets are captive ones. In the Free Software and Open Source movements we call captive markets "walled gardens" or "silos". But to most producers in the developed world, these are ideal. And to most consumers, they are business as usual.

Even after the Net obsoleted closed on-line systems, Yahoo, AOL and Microsoft continued to silo instant messaging inside their own walled gardens. In 2006, there should be no excuse for this.

Yet there is. We continue to believe, as both producers and consumers, that silos are okay. And worse, that a "free" marketplace is one where you get to choose the best silo.

We see this in the US today with our "choice" of services from phone and cable carriers. We even think the Net itself is a grace of telecom and cablecom carriage. After all, those are the guys we pay to get it. Those are the guys who have gradually increased our connection speeds.

Even the most broad-minded techies can get trapped inside the conceptual silo constructed by the telecom and cablecom carriers. As I write this, the debate over "Net Neutrality" is conducted almost entirely inside that silo.

The carriers claim to be fighting government regulation, when in fact they have known life only inside a regulatory habitat they built themselves and continue to control through an exceptionally powerful lobbying apparatus. Together with the lawmakers and regulators they control, the carriers have created what Bob Frankston (a father of both the spreadsheet and home networking) calls the Regulatorium.

The Regulatorium provides the building codes for telecom and cablecom silos. Telecom (including cablecom) "reform" is entirely about changing the building codes to make the silos more competitive with each other—not to free the captives of those silos or to blow the silos up altogether.

To the Regulatorium, a "free market" for Internet service means you get to choose between a cable and a telephone provider. That's it. These carriers can no more appreciate a truly free market than an agent in The Matrix can imagine a world not run by machines.

Carriers naturally oppose "Net Neutrality" because they've never provided it and don't want to start. No household customer of cable or ADSL service has ever experienced symmetrical, uncrippled service. For most of those same customers, the Internet is a secondary service offered alongside telephone or cable-TV services. As with human batteries inside The Matrix, a choice between two captors is all these customers have ever known. (Mobile telephony in the US is a little better, because there are more competing silos. But they're still silos. And not very open.)

On the other side of the debate are techies led by Google, Yahoo, Microsoft and other "content providers", all fighting

Accelerate your competitive advantage

Choose ServersDirect Systems powered by the innovative technology of the 64-bit Intel® Xeon® Processor with dual-processor functionality. Enjoy excellent performance and headroom for today's 32-bit applications. And protect your investments as you transition to 64-bit computing.

SDR-1500T

Designed for optimal performance while supporting maximum frequency dual-core Xeon® 5000/5100 sequence processors in a high-density 1U form-factor



\$2,499

Features

- 1U Rackmount Chassis with 700W Power Supply
- Intel® 64-bit Xeon™ 3.0/2x2M/667FSB (Dual Processor Option)
- Intel® 5000P (Blackford) Chipset
- Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2pcs x 512MB)
- 4pcs x WD4000YR, SATA 7200RPM hard drive
- 4 x 1" Hot-swap SATA Drive Bays
- RAID 0, 1, 5, 10 support
- Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller

2U Entry Level Server SDR-2500T

Optimize performance with dual-core Xeon® 5000/5100 sequence processors and FB-DIMM memory



- Intel® 64-bit Xeon™ 3.0/2x2M/1066FSB (Dual Processor Option)
- 2U rack designed with 550W Power Supply
- Intel® 5000P (Blackford) Chipset
- Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2x512MB)
- 1pc x Western Digital WD4000YR, SATA 7200RPM HD
- 6 x 1" Hot-swap SATA Drive Bays
- ATI ES1000 Graphics with 16MB video memory
- Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
- RAID 0, 1, 5, 10 support

\$1,999

3U Application Server SDR-3500T

64-bit + dual-core = power efficiency. This equation is ideal for intense computing environments and business-critical applications



- Intel® 64-bit Xeon™ 3.2/2x2M/1066FSB (Dual Processor Option)
- 3U Chassis with 650W Redundant Power Supply
- Intel® 5000P (Blackford) Chipset
- Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2x512MB)
- 1pc x 3Ware 9550SX-12 port RAID Controller Card
- 12pcs x Western Digital 250GB SATA RAID Drive
- ATI ES1000 Graphics with 16MB video memory
- Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
- RAID 0, 1, 5, 10 support

\$4,799

3U Database Server SDR-3500T

Outstanding performance with dual-core Xeon® 5000/5100. Ideal solution for storage and business applications.



- Intel® 64-bit Xeon™ 3.2/2x2M/1066FSB (Dual Processor Option)
- 3U Chassis with 650W Redundant Power Supply
- Intel® 5000P (Blackford) Chipset
- Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2x512MB)
- 2pcs x 3Ware 9550SX-8 port
- 16pcs x Seagate 300GB SATA-II Drive
- ATI ES1000 Graphics with 16MB video memory
- Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
- RAID 0, 1, 5, 10 support

\$5,899

5U Advanced Storage Server SDR-5500T

Powered by the latest dual-core Xeon 5000/5100® sequence processors, this system offers the best storage capacity available in a 5U format



- Intel® 64-bit Xeon™ 3.2/2x2M/1066FSB (Dual Processor Option)
- 5U Chassis, 24 hot-swap bays & 950W redundant power supply
- Intel® 5000P (Blackford) Chipset
- Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2x512MB)
- 1pc x 3Ware 9550SX-12 port RAID Controller Card
- 12pcs x Western Digital 400GB SATA RAID Drive
- ATI ES1000 Graphics with 16MB video memory
- Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
- RAID 0, 1, 5, 10 support

\$6,799

CONTACT US TODAY!

Our flexible online configurators allow you to create custom solutions. If you prefer, call us and our expert staff will help you assemble the perfect system. Servers Direct - your source for scalable, cost effective server solutions.

1.877.727.7887 | WWW.SERVERSDIRECT.COM



for Net Neutrality inside the Regulatorium: the carriers' home turf. The techies want the silo maintainers to play fair with Net traffic passing to captive customers inside the silos.

Regardless of how one defines "neutrality", locating the argument inside the Regulatorium requires admitting that the Net is a grace of silo'd carriage. Even pleas to restore "common carriage" (an enforced indifference by carriers to the contents or passengers they carry) serve to reify the silo'd nature of the carriers. In effect, these pleas say, "Because we are captive, and because you have no motivation to treat us all fairly, we must insist upon regulations that force you to do that." (Not surprisingly, this appeals to Democrats while offending Republicans, guaranteeing defeat in a Republican-majority Congress.)

One recent evening, I was explaining Net Neutrality to my wife, giving her a rundown of the various combatants in the debate and how they were doing. She listened patiently and replied on a level so high it reduced all the silos below to the dimensions of molecules. She said, "We're in the middle of a 100-year transition from analog to digital technology. That means we'll have another fifty years of prosperity and growth."

I realized instantly that the Net is not just about TCP or Neutrality or peer-to-peer or end-to-end or anything other than connecting digital devices across distances. It's about reducing those distances to zero—or as close to zero as possible. "The Internet is just a path", Bob Frankston says. And carriers, so far, have existed to create "billing events" in the middle of that path.

In "0 to 1 in 100 Years" (www.linuxjournal.com/node/1000056), an essay I wrote for the *Linux Journal* Web site in July 2006, I said this:

"Broadband" is like "long distance": just another name for transient scarcity. We want our Net to be as fast, accessible and unrestricted as a hard drive. (And in time even that analogy will seem too slow.)

The only way that will happen is if the Net becomes ubiquitous infrastructure—something which, in a practical sense, nobody owns, everybody can use and anybody can improve (worldofends.com/#BM_8).

There is infinitely more business in making that happen, and using the results, than Congress can ever protect for the carriers alone.

And guess who is in the best position to make money doing that? Right: the carriers.

Will somebody please tell them?

Of course that's a red herring. Telling something like that to a carrier is like telling an android to get a real body. Carriers can't do it. They're inside, not outside. They're a bus with a billing system.

You have to be free to see how absurd silos can be. You have to see markets as wide-open spaces opened by ubiquitous relationships, and potential relationships, between digital devices and the human beings who use them. You have to see unrestricted possibilities for the people and organizations putting those devices, their applications and their data to work. Those possibilities lose their limits once you set your

mind free of the notion that a free market is just a choice of silos.

So. What does it mean to be free?

Richard M. Stallman and the Free Software Foundation have something to say, at least in respect to software:

Free software is a matter of the users' freedom to run, copy, distribute, study, change and improve the software. More precisely, it refers to four kinds of freedom, for the users of the software:


- The freedom to run the program, for any purpose (freedom 0).
- The freedom to study how the program works, and adapt it to your needs (freedom 1). Access to the source code is a precondition for this.
- The freedom to redistribute copies so you can help your neighbor (freedom 2).
- The freedom to improve the program, and release your improvements to the public, so that the whole community benefits (freedom 3). Access to the source code is a precondition for this.

The cognitive linguist George Lakoff has come out with a new book, *Whose Freedom? The Battle Over America's Most Important Idea*. In a chapter titled "Why Freedom is Visceral", he says:

Freedom requires access—to a location, to an object, or to the space to perform an action. Access is a crucial idea in human thought...Freedom requires not just the absence of impediments to motion, but also the presence of access. Inhibiting freedom is, metaphorically, not just throwing up roadblocks, holding one back, taking away power, imposing burdens or threats or harm, but also failing to provide access. Freedom thus may require creating access, which may involve building.

I've said countless times that Linux (and the whole LAMP stack, which now potentially includes more than a hundred thousand components) is just building material. Its primary value is not in itself, but in what it makes possible—in what I call the "because effect". Far more money is made *because* of Linux than *with* it. Same goes for all kinds of other ordinary things, including buildings, cars and cell phones. The difference is that Linux was created in the first place with freedom as a primary purpose. So are countless other examples of free and open-source software. So is the Net.

"Free software is about free speech, not free beer", the Free Software Foundation says. In the same way, free markets are about free enterprise. In a free market, participants have the freedom to make and sell and service whatever they want. Free and open-source building components are commodities by intent. The abundance of a commodity may drive down prices for itself, but it also drives up the number and variety of purposes to which it can be put, and therefore drives up the number and variety of businesses it makes possible. When a commodity is free as in beer as well as free as in speech, the sum of business it makes possible is infinite. Yes, some of those businesses may be silos. But after a



Boost Performance on Linux Clusters.

Get your applications ready for scalable, parallel processing. Intel® Software Development Products help you create, analyze, and optimize high performance applications on Intel® processor-based cluster systems.

Intel® C++ and Fortran Compilers for Linux*

With features such as OpenMP* and auto-parallelism, the compilers provide highly-optimized threading features that help your applications take advantage of the performance benefits of Intel® multi-core processors.

Intel® Cluster Toolkit for Linux

Efficiently develop and distribute parallel applications with this multi-component software package that contains core libraries and analysis tools:

- **Intel® MPI Library**

Develop MPI-2 standard applications for all major cluster configurations and network architectures with one single high performance MPI library that now detects the optimal interface automatically.

- **Intel® Math Kernel Library Cluster Edition**

Add highly optimized, thread safe math routines such as BLAS, LAPACK, ScaLAPACK, FFTs, Sparse Solvers, Vector Math Library, and Random Number Generators to your applications.

- **Intel® Trace Analyzer and Collector**

Understand complex parallel program behavior by finding hotspots and bottlenecks in your distributed applications.

- **Intel® MPI Benchmarks**

Measure performance with this standard set of MPI benchmarks, formerly Pallas MPI Benchmarks (PMB)

“Intel’s MPI and Cluster Tools provide us the best cluster development environment. Using Intel Trace Analyzer and Collector, we were able to shorten MPI communication time by half by finding and removing bottlenecks with non-blocking and blocking communication patterns.”

*Dr. Takahiro Koichi
Computational Astro Physics Laboratory
RIKEN, Japan*

Put our products to the test! Download evaluation copies of the Intel® Software Development Products today at www.intel.com/software/products.

Find a reseller at www.intel.com/software/products/reseller.

Come visit us at SC06 in booth 1523.



NEW!


TS-7300 High Security Linux FPGA Computer




\$219 qty 1 **\$189** qty 100

200 MHz CPU

- User-programmable Altera 2C8 Cyclone II FPGA
- PC/104 expansion bus
- Fanless, no heat sink
- 10 RS232 serial ports
- 55 DIO ports
- 2 10/100 Ethernet
- VGA video out
- matrix keypad & text mode LCD
- 2 USB ports
- 2 hot-swappable SDCard sockets
- SD is software RAID1 capable
- Linux, NetBSD
- Real Time extension

 Design your solution with one of our engineers

- Over 20 years in business
- Never discontinued a product
- Engineers on Tech Support
- 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



We use our stuff.

visit our TS-7200 powered website at

www.embeddedARM.com

(480) 837-5200

COLUMNS

LINUX FOR SUITS

while, it becomes clear that the cost of maintaining silos is exceeded by the benefits of freedom—for everybody involved.

In *The Matrix Reloaded*, the hero of the story, Neo, meets the architect of The Matrix. With the hauteur and condescension of a stuffy college professor, the architect leads Neo to a conclusion. “Choice”, Neo says. “The problem is choice.” The architect agrees, and presents Neo with two choices, each leading to a determined outcome. Neo doesn’t agree. He sees more possibilities than the architect, because Neo is human. He is not a program. He is capable of creation at a level machines cannot achieve, no matter how intelligent they may be. Human beings can be profoundly creative, and vastly original.

With original technologies, we can be exceptionally good at opening possibilities, at multiplying choices. This is what we do with free and open-source software.

What we do *with* Linux is also what we want *from* it. Both are about choice. You can’t get full openness to possibilities—and choices among them—without full access to source code, and the freedom to improve that code as well as put it to use. In the techie world, we know about the virtues involved, and the effects as well.

But there is another path of access to sources that needs to be opened. This is the one between users and engineers—between the people who use code and the people who write it. The same goes for makers of hardware and non-software products. In fact, just about everything. Yes, there are trade secrets and recipes and drug formulas and other kinds of stuff for which one can justify isolating creators from users. But in most cases, creators and users can benefit far more from interaction than from isolation.

In the free world we’re building, isolation of creators is also becoming increasingly impractical. What David Weinberger called “Fort Business”—the building with the working space isolated inside physical as well as digital firewalls; the badges for employees, the escorts for visitors, the nondisclosure agreements for contractors—is becoming more anachronistic and absurd by the minute. Workers operate in the outside world as well as the inside one. Information customarily kept secret may prove to be more valuable in the open.

It is still common in Silicon Valley to see, in the vast empty cubicle farms of failed companies, posted warnings not to let secrets out of the building. This is the kind of insanity we have come to expect from marketing: a category of activity that works to control customers on the outside and creativity on the inside.

Marketing, advertising and PR have lots of warm BS about the good that they do. But the verbs give away their true intentions. Marketers want to “grab”, “capture” and “hold” people’s attention. Customer “Relationship”

Management (CRM) systems are mostly devoted to Silo Inhabitant Retention.

Yes, I know that’s wrong and unfair and not even close to marketing’s ideals. But there is one good reason why customers hate advertising, journalists hate PR and engineers hate marketing. Put simply, it gets in the way.

In most companies, marketing stands between users and engineers. Obstruction may not be formal, but it’s there. We have this notion—programmed, Matrix-like, into Business As Usual—that says we need marketing to guide creation on the inside and to produce demand on the outside. We envision a process of product (or service) creation and sales that works in a linear, value-chain way. Feedback from the market comprises the return side of a cycle that needs to be formalized and run through a System. Marketing is in charge of that. Engineers aren’t. They’re back at the beginning of a process that doesn’t involve them directly with customers or users. That’s up to sales, marketing, customer support or some other non-engineering part of the corporate bureaucracy.

That whole mentality is a silo too. Maybe even a Matrix. Because we’re trapped inside it, and we need to free our minds from it.

Survival and prosperity for technical companies in the long run will require more and more contact between makers and users. As with the trapped inhabitants of The Matrix, we need to start over in a new world that lacks the familiar comforts of the old one. There was comfort in the distance put by the silo between engineer and user. The engineer could work in peace.

But that distance is gone. That peace is gone. We have to make a new world where we know first-hand how our stuff works for real customers and real users. We need to engage and relate. Reports from marketing on studies and focus groups aren’t going to cut it anymore. We’ll need to know more and more on a first-hand basis.

Of course, engineers will still enjoy their privileged positions. “Show me the code” will always be the base price of admission to the software meritocracy. But meritocracies will be built on meritocracies. More products, more services, more possibilities, more choices, more opportunities, more value and more money. For everyone involved.


Code comes down to ones and zeros. We’re building a whole new world with those. My wife is right. It will take a long time. We’ll make a lot more business happen—faster—by opening markets up than by closing them down. And we’ll get a lot more done by connecting freedoms to create with freedoms to choose and buy and use and re-use.

Liberating ourselves won’t be easy. But it’s necessary. ■

Doc Sealts is Senior Editor of *Linux Journal*.

AWARD WINNING SERVERS AND STORAGE

ABERDEEN STIRLING 54T



1U 14 Inch Depth Mini Rack

- Intel® Pentium® D processor powered by Intel dual-core technology with up to 1066FSB and 4MB Cache
- Intel E7230 "Mukilteo" Chipset with 64-Bit Support
- Up to 8GB DDR2 667/533MHz Memory
- Up to 750GB SATA Hard Drive
- 260W AC Power Supply w/PFC
- 5-Year Warranty

Starting at **\$999**

ABERDEEN STIRLING 106T




1U Short Depth 4 SATA 3TB

- Intel® Pentium® D processor powered by Intel dual-core technology with up to 1066FSB and 4MB Cache
- Intel E7230 "Mukilteo" Chipset with 64-Bit Support
- Up to 8GB DDR2 667/533MHz Memory
- Up to 4 x 750GB (3.0TB) Hot-Swap SATA Hard Drives
- 300W AC Power Supply w/PFC
- 5-Year Warranty

Starting at **\$1,169**

ABERDEEN STIRLING 128T



1U Dual-Core 4 SATA/SCSI 3TB

- Up to two Dual-Core Intel® Xeon® processors with up to 1333FSB
- Intel® 5000P "Blackford" Chipset with 64-Bit Support
- Up to 64GB 667/533MHz Fully Buffered ECC DDR2 SDRAM Memory
- Up to 4 x 750GB (3.0TB) Hot-Swap SATA Hard Drives
- Up to 4 Hot-Swap SCSI Hard Drives version available (Stirling 128S)
- 700W AC Power Supply w/PFC
- 5-Year Warranty

Starting at **\$1,999**

ABERDEEN STIRLING 226T



2U Dual-Core 6 SATA/SCSI 4.5TB

- Up to two Dual-Core Intel® Xeon® processors with up to 1333FSB
- Intel® 5000P "Blackford" Chipset with 64-Bit Support
- Up to 32GB 667/533MHz Fully Buffered ECC DDR2 SDRAM Memory
- Up to 6 x 750GB (4.5TB) Hot-Swap SATA Hard Drives
- Up to 6 Hot-Swap SCSI Hard Drives version available (Stirling 226S)
- 550W AC Power Supply w/PFC
- 5-Year Warranty

Starting at **\$2,125**

ABERDEEN STIRLING 228T



2U Dual-Core 8 SATA/SCSI/SAS 6TB

- Up to two Dual-Core Intel® Xeon® processors with up to 1333FSB
- Intel® 5000P "Blackford" Chipset with 64-Bit Support
- Up to 64GB 667/533MHz Fully Buffered ECC DDR2 SDRAM Memory
- Up to 8 x 750GB (6.0TB) Hot-Swap SATA Hard Drives
- Up to 8 SCSI/SAS Hard Drives version available (Stirling 228S/228Q)
- 700W High-efficiency Redundant Power Supply
- 5-Year Warranty

Starting at **\$2,485**

ABERDEEN STIRLING 328S



3U Dual-Core 8 SCSI 2.4TB

- Up to two Dual-Core Intel® Xeon® processors with up to 1333FSB
- Intel® 5000P "Blackford" Chipset with 64-Bit Support
- Up to 32GB 667/533MHz Fully Buffered ECC DDR2 SDRAM Memory
- Up to 8 x 300GB (2.4TB) Hot-Swap SCSI Hard Drives
- 2 x 5.25" Peripheral Drive Bays
- 550W Fixed or 760W Triple Redundant Power Supply
- 5-Year Warranty

Starting at **\$2,445**

ABERDEEN STIRLING 428T



4U/Tower Dual-Core 8 SATA/SCSI 6TB

- Up to two Dual-Core Intel® Xeon® processors with up to 1333FSB
- Intel® 5000P "Blackford" Chipset with 64-Bit Support
- Up to 64GB 667/533MHz Fully Buffered ECC DDR2 SDRAM Memory
- Up to 8 x 750GB (6.0TB) Hot-Swap SATA Hard Drives
- Up to 8 Hot-Swap SCSI Hard Drives version available (Stirling 428S)
- 800W High-efficiency Redundant Power Supply
- 5-Year Warranty

Starting at **\$2,485**

ABERDEEN STIRLING X633



6U Dual-Core 32 SATA 24TB

- Up to two Dual-Core Intel® Xeon® processors with up to 1333FSB
- Intel® 5000P "Blackford" Chipset with 64-Bit Support
- Up to 32GB 667/533MHz Fully Buffered ECC DDR2 SDRAM Memory
- Up to 32 x 750GB (24.0TB) Hot-Swap SATA Hard Drives
- 2 x Rear Hot Swap SATA Hard Drives for OS
- 1350W Triple Redundant Power Supply
- 5-Year Warranty

Starting at **\$7,999**

"The Ultimate Linux Server... too fast for our benchmarks... we recommend the Aberdeen line of servers without reservation."

Linux Journal, August 2006

"powerful enough to tackle the most cutting-edge applications"

CRN Test Center Recommended, August 2005

"staggering ... eye-opening ... the highest WebBench numbers to date"

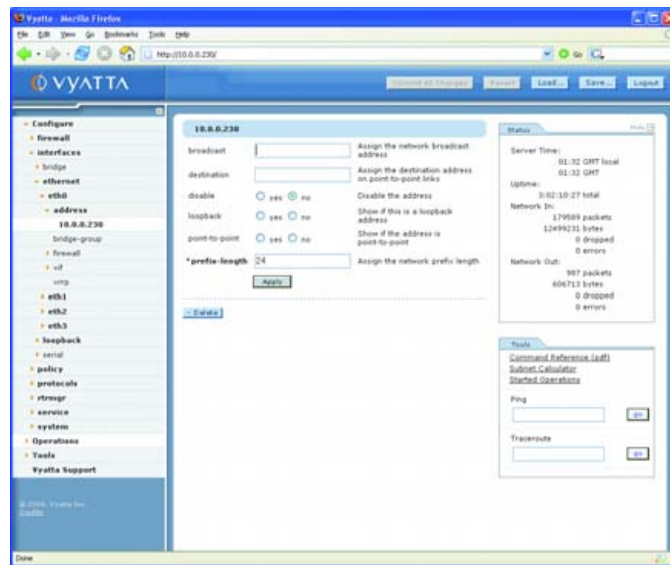
PC Magazine, February 2005



Vyatta's Open Flexible Router

The word vyatta means "open" in ancient Sanskrit, and this relatively new company that has taken on the moniker is applying this philosophy of openness to routers. Vyatta's product, the Open Flexible Router (OFR), is reputed to be the IT industry's first enterprise-grade, open-source router platform. OFR allows users to utilize industry-standard x86 hardware to "create a high-performance router for WAN and LAN routing" that offers "dramatically improved price/performance and open flexibility" vis-à-vis closed-source solutions. Vyatta targets OFR at mid-sized enterprises or branch offices of larger ones. Product benefits include not only all standard routing protocols and high-availability and security features, but also the ability to customize the product and add features as needed. The latter is intended to give users flexibility in managing future requirements on their own terms rather than relying on the actions of closed-source vendors. In addition, both free and paid support options exist. Free support comes in the form of the Vyatta Community, which includes tools, discussions, blogs and newsfeeds. Paid support includes tiered subscription packages customized to the user's needs. The OFR software is available for free from Vyatta's Web site.

www.vyatta.com



SageTV's Media Center for Linux

Hear ye, hear ye, all TV freaks! SageTV recently released a product dubbed SageTV Media Center, which, according to the the company, allows you to turn a Linux-based computer into a full-featured PVR and media center. In addition, you can utilize the Placeshifter option to watch your entire home-based media library from any remote computer with high-speed Internet access. Furthermore, the Media Center also can interact with SageTV's related product, the Wireless Media Extender, which allows every TV in the home to access all live or recorded media content from the Media Center independently without a need for a PVR on each set. Meanwhile, a central, unified media library is maintained. SageTV claims that one can set up a complete, full-home media center for hundreds rather than thousands of dollars for other solutions, and there is no need to pay subscription fees for PVR capabilities.

www.sagetv.com

Sage Software's Sage Accpac ERP

Sage Software (no relation to SageTV above) recently announced a new version, now Release 5.4, of its Sage Accpac ERP business management system. In essence, Sage Accpac ERP is a Web-enabled accounting and business management solution for integrated business management applications, including CRM, POS, HR, inventory management and so forth. It can be deployed either on-site or on-line. New user features in Version 5.4 include a new returns management (RMA) module and 267 other product enhancements, such as improved multi-currency tax accounting. A new technical feature of Accpac ERP is its compatibility with Intel and AMD 64-bit processors. Also of interest to the Linux crowd is the ability of Sage Accpac ERP to run simultaneously on both Linux and Windows clients in parallel against the same database, be it DB2, Oracle, PSQL or Microsoft SQL.

www.sageaccpac.com

Sales Optimizer Summary			
Results			
	Yr to Sep 01	Yr to Sep 00	
Revenue	4,364,327	4,453,638	●
Cost of Goods	2,370,776	2,362,173	●
Margin	1,993,551	2,091,465	●
Units	306,835	312,191	●
Margin %	45.68%	46.96%	●
Activity			
Product	240	237	●
Customer	239	239	●
Salespersons	11	11	●
Transactions	20,195	20,222	●
Best of Breed			
	Revenue	Margin	Units
Customer	Revenue	Margin	Units
WB Market LA	484,634.88	226,276.26	32,422
WB Market Las Vegas	338,740.53	160,267.83	22,861
TL Food Market	243,367.14	113,842.31	17,031
Excel Foods - Boise	220,297.60	113,296.23	14,078
Ez Ze Stop	176,502.40	86,091.64	11,312
TM Market	117,081.08	59,752.92	8,172

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



Kerio's MailServer

The Kerio folks told us that they have updated their MailServer product, now Release 6.2, which supports Linux, Mac OS and Windows. The product is a groupware server targeted to small- and mid-market businesses and is available in 16 different languages. The most noteworthy new user and administrator improvements are Universal binaries (MailServer is now a Universal application on Mac OS X); two-level spam rating (messages can be tagged or blocked/rerouted); distribution-list definition and syncing; Sticky Note support (users can create and sync Sticky Notes in WebMail and Outlook); Clam AntiVirus plugin; and improved search, including calendar events and tasks. Kerio also improved its WebMail application utilizing Ajax.

www.kerio.com

Learning Maya 7: Foundation with DVD by Autodesk Maya Press (Sybex)



An Imprint of WILEY

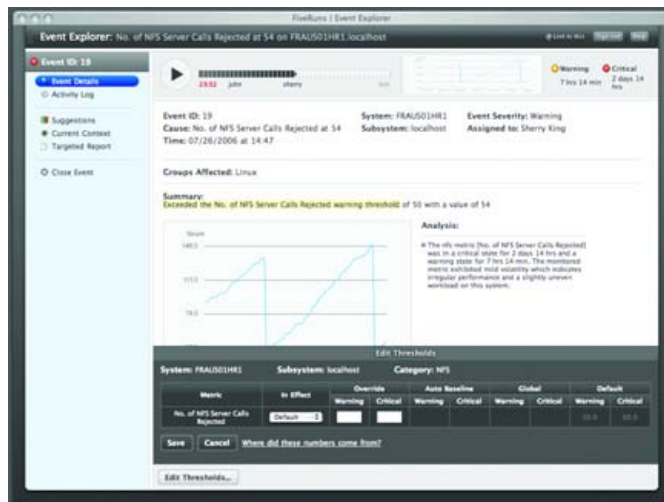
Those of you who follow this section will wonder why products in the high-end 3-D graphics space keep popping up. Admittedly, your editor's dream to work for Pixar has (foolishly) not yet been dashed. This month, we'll see a new book title about Autodesk's Maya package, titled *Learning Maya 7: Foundation with DVD*. The book is authored by Maya's creators at Autodesk and published by Sybex's Autodesk Maya Press imprint. Maya is a powerful program used by the film, TV and computer/video game industries for 3-D modeling, animation, effects and rendering. And, Maya runs on Linux as well. This full-color book uses a real, forthcoming Hollywood film as material for its tutorials, teaching readers animation, modeling, texturing and visual effects. Also included is a DVD with instructor-led tutorial videos and other extra features. (Psst! Anyone out there with connections at Pixar who can hook me up?)

www.sybex.com

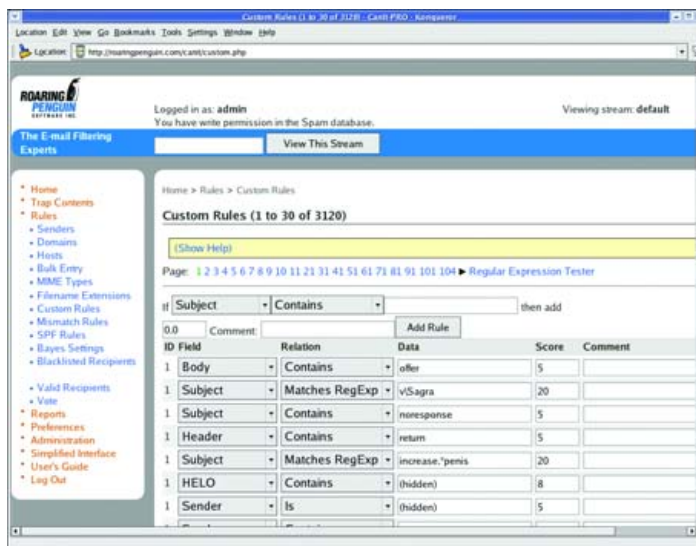
FiveRuns' Systems Management

FiveRuns, whose name comes from the founder's favorite childhood fishing hole, has left the world of beta releases (now 1.0) with its open-source Systems Management product, "complete with Web 2.0 richness and usability." According to FiveRuns, its "approach to systems management utilizes open-source technologies, such as Ajax, Ruby and Ruby on Rails to deliver a sleek, hosted application." Because IT professionals must typically manage both open-source and proprietary environments, FiveRuns' philosophy is to support these heterogeneous environments through a single interface. FiveRuns also claims that businesses of all sizes can benefit from its products. Other features include software-as-a-service-based monitoring, allowing users the flexibility to subscribe to the appropriate level of service as infrastructure levels change, as well as automated configuration discovery and auto-baseline thresholds to keep implementation time to a minimum. Free support for the product is available through FiveRuns.net, a community of users that includes troubleshooting advice. A free 30-day trial of FiveRuns is available by signing up on the company's Web site.

www.fiveruns.com



Roaring Penguin Software's CanIt-Domain-Pro



The developers of e-mail applications are rockin' these days, and one of the companies responsible for the buzz is Roaring Penguin Software, who just released Version 1.0 of its CanIt-Domain-Pro. This product is an anti-spam solution designed specifically for organizations with complex e-mail needs related to managing multiple domains, such as ISPs, universities, governments, Web hosts and so on. Roaring Penguin says its product can solve problems in universities and governments

"where each department often has its own domains" or in the ISP world where "merger activity has brought multiple domains under one roof". Roaring Penguin's approach is to simplify the management of multiple domains by allowing a "super system administrator to delegate system administration rights" to lower-level "realms", with realm administrators able to create baseline rules, users and streams within their jurisdictions. CanIt-Domain-Pro is available as a software download, a soft appliance or a preconfigured plug-and-play appliance.

www.roaringpenguin.com

Call MisterHouse TO REGULATE YOUR HEAT

Combine MisterHouse with Perl scripts, X10 and the Linux kernel to create an automated zoned radiant heat system.

DAVID LYNCH

I am an architect and a software consultant, I am building my own home, and both my wife and I have allergies. Forced-air environmental control systems are not good for people with allergies. I have chosen to use radiant heating. Environmental control systems are one of the areas of home automation into which hobbyists rarely venture. Although the control systems for air conditioners and heat pumps are not particularly complex, a poor control system can wear out an expensive compressor early, if it does not destroy it. This taboo against directly controlling environmental control systems hardware extends to

systems that are neither as easy to destroy nor as costly to replace.

This article discusses controlling a home radiant heating system using a home control system running under Linux. I have used several guiding principles in designing and implementing this system. I am a huge proponent of KISS (Keep It Simple), and I am an underutilized software consultant from a community that still believes pennies are money—I am cheap. I think that is a virtue. Simple trumps complex, cheap prevails over costly, software substitutes for hardware, do-it-yourself replaces buying or hiring.



ILLUSTRATION: ©iStockphoto.com/Bryce Kroll

Division into Zones

Radiant heating systems are not particularly complex. My house is divided into zones, and each zone has its own controls. To some extent, the control system for the zone interacts with controls for the whole system. In a radiant heating system, when a zone requests heat, water is heated and circulated into something to radiate that heat into the rooms making up that zone.

System Components

The first system component is a heat source, and there is a wide variety of possible heating sources. I am using a 144,000 BTU wall-mounted instantaneous natural gas hot water heater with a 95% efficiency rating. One of the features of these units is that they are regulated and controlled internally. Simply drawing water triggers the electronic ignition. The unit itself controls its output to match demand. My unit has some ability to manipulate the preset water temperature as well as monitor its performance. Regardless of whether it is conventionally controlled or controlled by a home automation system, my water heater itself requires no external controls.

The next component is a heat sink—the way heat gets from the distribution system into the zone. In my house, this is PEX tubing, approximately 16" on center inside a concrete slab. There are many other ways of getting heat from the water into the zone, including baseboards, radiators and tubing stapled to the underside of floors. Each method has its own properties, but the most critical of these is the amount of heat it can radiate into the zone. PEX tubing in a concrete slab is one of the more efficient means of transferring heat to a zone. The maximum radiated heat from a concrete floor with radiant tubing is between 50–90K BTUH per 1,000 square feet or 50–90K BTUH for 1,000 square foot of floor—depending on how warm a floor you are willing to tolerate.

Proper design of the heat source and sink are beyond the scope of this article; however, they are not really complex. The heat sink must be

capable of supplying more heat than the zone can lose in an hour under the most extreme conditions. The heat source must be able to provide at least as much heat as all zones could demand under the most extreme conditions. Under cold but less than extreme conditions, I do not want the system running most of the time; the norm is for the heat to run about one-third of the time. All heating systems are over-designed.

Another component is the circulator(s). Water must be persuaded to flow through the system. The norm is that at least one pump or circulator is required. Circulation systems can be designed in a number of ways. I have been able to procure small circulators very cost effectively (cheaply), so I have used a circulator for each zone. This significantly simplifies many other aspects of the design of the system. It eliminates the need for zone control solenoids and their controls. It eliminates controls for the circulator separate from those of the zone. It makes providing a zone-by-zone failsafe override easier, and it reduces the need to balance the system.

Another part of the system is a closed-loop water distribution system. Hot water needs to get from the heat source to the heat sink and back. In my home, this is more PEX tubing, as well as manually constructed manifolds.

Finally, there is the control system, and that is where Linux comes in.

Premises

There must be a failsafe at least to keep the pipes from freezing should the home automation system fail. I have no provisions for dealing with an extended loss of power during severely cold weather—very few modern heating systems will work without power. Even the failsafe will fail absent another source of power to drive the circulators and ignite the water heater in the event of a prolonged power failure. A failsafe is accomplished by wiring a traditional cheap thermostat in parallel with the home control system on a sufficient number of dispersed zones to maintain temperatures above freezing.

Feedback is more important than reliability—it is not important that a



Figure 1. Takagi Tk Jr 144K BTUH Hot Water Heater



Figure 2. Radiant Distribution Manifold and Circulators

zone starts heating at a specific instant. It is important that over the course of time, the zone temperature does not wander too far outside the tolerance from the current setpoint for that zone. That tolerance is critical. Expecting to maintain a temperature of exactly 70°F in a zone all the time is unrealistic. Doing so would make the system turn on and off constantly. All environmental control systems have some tolerance built in. The primary difference in my system is that both the setpoint and the tolerance are controllable. Feedback is important in that it allowed me to use an extremely inexpensive means of controlling AC power: X10 appliance modules.

X10 is fairly ubiquitous. It is readily available and cheap. In a lot of installations, it is highly reliable, but some are highly troublesome. Feedback is also important because if all else fails, my heating system needs to attract my attention to solve the problem.

The controls and sensors need to be cost-effective devices easily interfaced to and managed by the home control computer—did I mention that I am cheap?

The Control System

The X10 appliance modules to turn on and off AC power to the AC devices (circulators) that are being controlled are Dallas Semiconductor DS1820 1-Wire temperature sensors. These are extremely inexpensive and are an easily interfaced means to sense temperatures. Each zone may have one or more temperature sensors. Eventually, I intend to place sensors at different heights within the zones to explore the radiant heating premise that rooms with warm floors are perceived as warmer even if the room temperature higher up is cooler. Furthermore, I monitor the supply and return temperatures for each zone. This provides feedback, and it provides a means of monitoring system performance. The heat supplied to each zone is the differential between the supply and return temperature, the flow rate and the cycle time. The flow rate for the circulators is approximately constant, and all other factors can be monitored. This means the output and performance can be monitored in real time.

Finally, there is a need to determine human input—I need to be able to demand that my living room be hotter or colder. This is done in a number of ways.

For my first heating season, a single \$10 US thermostat was used to turn on all circulators for all zones concurrently. At initial startup, the hot water heater started about a minute after the circulators and ran for about 17 hours straight. I was starting to panic, wondering how large a loan I would need to afford heat. Then it stopped. After that it became a game trying to catch the system running the rest of the winter. Once each zone was up to temperature, they very rarely ran.

During the second season, I used a simple MisterHouse Perl script. Changing a zone's setpoint involved simple changes to that Perl script. It worked, but it was not a human-friendly user interface.

Ordinary thermostats can be installed and connected to digital inputs. There are numerous ways to do this, and there are relatively cheap 1-Wire digital input sensors. The thermostats are not used to set the temperature; they are used as inputs to determine whether I want the temperature altered. I use these primarily because all users, regardless of their competence with computers, tend to be comfortable with and understand them, which is important in a guest room. However, because they are inputs to the control system, they constitute requests, not demands, and they are not a failsafe.

The user interface for the normal computer-savvy users (my wife and I) is an evolving Web interface with status and alterable settings for each zone. The Web interface can be used from any computer on the home net, from a wireless PDA supporting a Web browser or even from remote locations, provided sufficient security precautions are taken.

Nothing about an environmental control system is particularly time-sensitive. I develop embedded and sometimes real-time systems. Environmental control is not real time. This means the demands of the control software on the Linux system are relatively low and not particularly time-critical. In fact, in my home, the heat source for the radiant heat is

shared by domestic hot water. This has advantages and disadvantages and is not permitted by some building codes. When domestic and heating demands exceed the capacity of the hot water heater, such as when someone is taking a hot shower, domestic needs trump those of heat. It is unlikely I will notice if the heat delays a few minutes before starting. But if I am in the shower, I will notice if I am deprived of hot water just so the heating cycle can start immediately.

The open-source MisterHouse Project provides a capable and highly programmable home control system, with capabilities well beyond those required for this project. It runs on Linux and Windows and supports a wide variety of controls, sensors and other hardware. It provides a very capable Web interface that I was able to extend easily to support my zone controls. And, it has features, such as floor plan integration, that I hope to take advantage of later. There already have been several environmental control systems implemented using MisterHouse. These have used expensive computer-monitorable thermostats controlling traditional HVAC control systems. In my implementation, MisterHouse does all the heavy lifting.

MisterHouse is written in Perl and has provisions to allow MisterHouse users to incorporate their own Perl routines easily, as well as an API with functions and events targeted at home control. It is fairly simple to create a bit of Perl code and have it execute every minute, every three minutes or only on Sundays when there is a full moon. It is easy to monitor or control 1-Wire and X10 devices—all the elements needed to make a working system.

I love programming in Perl, because it is easy, powerful and forgiving. You do not need to be a Perl monk to create custom scripts for MisterHouse. The MisterHouse distribution and Web site include numerous examples, many of which are only a few lines of fairly simple code. A substantial amount of home control can be accomplished with MisterHouse without doing any programming at all.

I primarily used Debian Linux, and I created a Linux Vserver (lightweight virtual server), specifically for MisterHouse. This is not strictly necessary, but it is cheaper and easier than a dedicated machine, and it's simpler and cleaner than running a mess of different daemons on a single machine. I highly recommend Vservers; they make experimenting with configurations fun and easy.

MisterHouse can be installed under Debian with apt:

```
apt-get install misterhouse
```

It requires a collection of Perl modules, and these dependencies should be taken care of by apt.

While planning this system, I decided on a computer X10 interface, the ACT TI-103, specifically because the most common X10 interfaces, the CM17 and CM11, have been known to have issues. The TI-103 sends a stronger than normal X10 signal and can receive a weaker than spec X10 signal. Furthermore, it supports X10 extended addressing, allowing the use of more than 255 X10 devices in a single home. Unfortunately, I mistakenly assumed that there would be MisterHouse support for it.

For early development, I used an old CM11 I had lying around, but eventually, I wrote a new MisterHouse driver for the ACT TI-103 that is now part of the MisterHouse Project. The ACT TI-103 is a nice controller and was easy enough to talk to—with some effort you can manipulate it directly using minicom as it uses ASCII command strings. Developing the TI-103 driver was more complex and time consuming than the rest of my MisterHouse HVAC software. So far, it has not proven to be any more or less reliable than the CM11. But, my TI-103 driver has been part of MisterHouse distributions since MisterHouse 1.102. After I was well underway with this effort, Insteon came out with a new series of power-line devices that are superior in many ways to X10 devices. However, they are still more expensive than X10 modules. Neil Cherry should have MisterHouse support for Insteon completed by the time this article is published.

MisterHouse needs to be configured for the controllers and sensors. The X10 controller is set by assigning the correct port, such as /dev/ttyS0, to the correct X10 device in the MisterHouse configuration file /etc/misterhouse/mh.ini:

YOU DO NOT NEED TO BE A PERL MONK TO CREATE CUSTOM SCRIPTS FOR MISTERHOUSE.

► FEATURE: REGULATE YOUR HEAT

```
cm11 = /dev/ttyS0
```

or:

```
ti103 = /dev/ttyS0
```

You may want to configure a number of other parameters, such as latitude and longitude, that would make MisterHouse more useful, or if you wanted to make heating decisions based on the weather report but nothing else specific to heating control.



Figure 3. Radio Shack/X10 Appliance Module Controlling Circulator

X10 devices are set up in a separate configuration file, `x10.mht`. The parameters for each device are the type of X10 module, its address, name and the groups to which it belongs. I needed an entry for each appliance module controlling the circulator for each zone. An example entry would be:

```
X10A, C1, Circulator1, All_Heat|Night_Setback
```

for an X10 appliance module with a C house code and a 1-unit code named `Circulator1` included in the groups `All_Heat` and `Night_Setback`. In a MisterHouse Perl script, one way this can be turned on is with the code set:

```
$Circulator1 'ON';
```

and all `Night_Setback` devices can be turned off with:

```
&set_all($Night_Setback, 'OFF');
```

If you have multiple X10 controllers, you may need to assign a controller to the device. And, it is very useful to assign labels and groups to devices. `mht` files are transformed into `mhp` (Perl code) files as part of MisterHouse startup,

or they can be reloaded on a running system through the MisterHouse user interface. Once X10 devices are correctly entered into the `.mht` file and MisterHouse starts up, they can be controlled manually by the UI, or they can be controlled programmatically via bits of Perl. It also is possible to add X10 devices to MisterHouse directly from the MisterHouse Web interface.

MisterHouse supports a number of user interfaces, including a Tk interface and Telnet. But, the primary interface I use is the built-in MisterHouse Web server. By default, it should come configured to work and should be on port 8080. Using a Vserver made it easy to move it to the standard http port 80. Inside my home, browsing mh.dlasys.net brings up the MisterHouse main page.



Figure 4. DS1820 1-Wire Test Rig

I had several available alternatives for interfacing MisterHouse to 1-Wire. MisterHouse can communicate directly with 1-Wire devices through a 1-Wire serial interface, and they are configured in an `.mht` file much like X10 devices. One of the machines I was planning on using did not have sufficient serial ports for both the X10 controller and the 1-Wire serial controller, so I opted for a 1-Wire USB controller. MisterHouse currently does not directly support the 1-Wire USB controller. But there are several Linux-specific means of using the 1-Wire USB controller. The Linux 2.6 kernel recently added support for a number of 1-Wire devices and interfaces. This is what I would have preferred to use. Unfortunately, at the time I was developing the code, it was very new and not well documented. There is a 1-Wire filesystem that maps 1-Wire devices to a filesystem. `digitemp`, a standalone program for both Linux and Windows, can read 1-Wire devices from both USB and serial interfaces, and that is what I elected to use.

My first implementation had `digitemp` polling all the DS1820 temperature sensors inside MisterHouse. My current implementation has `digitemp` polling the temperature sensors as an external cron job and recording the information into an RRD database. The following Perl script creates an RRD database with entries for each DS1820 in the `@sensor` list:

```
#!/usr/bin/perl
# Category = HVAC
@sensor = ( "28E8E30500000083", "2853327C000000D4" );

my $cmd = "rrdtool create temp.rrd --start N --step 60 ";
for $i ( 0 .. $#sensor ) {
    $cmd .= "DS:$sensor[$i]:GAUGE:120:30:180 ";
}
$cmd .= "RRA:MIN:0.5:60:8760 ";
$cmd .= "RRA:MAX:0.5:60:8760 ";
$cmd .= "RRA:AVERAGE:0.5:60:8760 ";
print "$cmd\n";
my $result = system($cmd);
```


Polywell Server and PC Solutions

Specialize for Small to Medium Size Business

High-End Servers and Workstations

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

Low-Cost Linux PCs

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

- 2 x AMD® Optron™ Dual-Core Processors 285+ with Hyper Transport Technology
- Upto 64GB 400MHz ECC DDR (16 Sockets)
- Upto 2TB 4 x 500G Swap Drive per 1U Rack
- 1 x PCI-X 133MHz or 1 x PCI-E RISER Slot
- Dual Gigabit Ethernet, ATI Graphics, 4 SATA-RAID
- CD-ROM Drive, Optional DVD-RW or CDRW
- Optional Slim Floppy Drive
- 1U 24" Depth Rack Chassis with upto 600W PS
- 4 x Swappable Drive Bays (SATA or SCSI)
- Supports Linux, FreeBSD or Windows
- Custom Configuration Available
- Please call for other Options



1114AIS-2050M, 285+,32GB, 2TB **\$9,999**

1114AIS-2050M, 260HE,16GB, 1TB **\$4,499**

1112ES-2200A, 265+, 2GB, 500GB **\$1,999**



4024AIS 4U 12TB SAS/SATA **\$12,500**

4024AIS 4U 6TB StorageServer **\$7,950**

Linux Appliance PCs

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

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



626TD112-800LX
OEM Appliance starts at **\$299**

64GB DDR 4-way Workstation

- 2 x AMD® Optron™ Dual-Core Processors 285+
- Upto 64GB ECC DDR (16 Memory Sockets)
- 2 x 133MHz PCI-X, 1 x PCI-E x16 Slots
- Dual Gigabit Ethernet, 4 x SATA-RAID Controller
- Quiet and Cool 11-Bay Tower + 1400W 80Plus P/S
- 2x300G HD, DVD-RW, Floppy, Optional Card Reader
- 8-Layer Motherboard with Special Quiet Cooling
- Optional Sound Card, 1394 firewire
- Supports 64/32-bit Linux, FreeBSD or Windows
- Special for Large Memory Intensive Applications
- Built-to-Order or Configure-to-Order

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



64G RAM, 2x285+, QuadroFX5500 **\$22,999**

32G RAM, 2x265+, QuadroFX3450 **\$7,599**

16G RAM, 2x246+, QuadroFX1500 **\$3,750**

PolyStation 2050M
4-Way Two Dual-Core Processors
Up to 64GB memory, 16 Sockets

1U Power Saving ISP Server

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



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

Pick the style for your application



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

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



5U 16-way 865, 32GB, 2TB, 8801T5U **\$19,999**

2U 8-way 865, 8GB, 900GB, 8422C **\$6,999**

Mainstream Linux Systems

- AMD® Athlon™64 or X2 Dual-Core Processors with Hyper Transport Technology
- 512MB to 4GB Dual Channel DDR Memory
- 80GB to 500GB Hard Drive (SATA-RAID-0/1/5)
- 2 PCI, 2 PCIe, Gigabit LAN, 7.1 Audio, 1394a
- nForce 430 Chipset, GeForce 6100 Graphics
- Desktop, Tower, MiniBox or Rackmount Chassis
- Internet Server or Linux Appliance Configuration

939NV3000TD12LX starts at **\$399**

Polywell OEM Services, Your Virtual Manufacture
Prototype Development with Linux/FreeBSD Supports
Small Scale to Mass Production Manufacturing
Fulfillment, Shipping and RMA Repairing



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

888.765.9686

www.Polywell.com/us/LJ



Polywell has been in business since 1987, Our professional engineers have many years of experiences helping small to mid-size business to custom make or migrate products into Linux Appliances. We also provide Linux/FreeBSD Driver Development Services.

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

Optron, 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.

▶ FEATURE: REGULATE YOUR HEAT

DS:\$sensor[\$i] specifies the data source (each DS1820), 30:180 are the min and max values for the data source and GAUGE is a particular RRD data source appropriate for temperature sensors. There are three archives: a Min, Max and Average.

I live in the US, so all temperatures are in Fahrenheit; however, digitemp can read and record either Centigrade or Fahrenheit values, and all other values throughout the system are relative to those reported by digitemp.

Polling is accomplished with the following script:

```
#!/usr/bin/perl
# Category = HVAC

@sensor = ( { sn => "28E8E30500000083", temp => 0 },
             { sn => "2853327C000000D4", temp => 0 } );

# Read the output from digitemp
open( DIGITEMP, "digitemp_DS2490 -c /etc/digitemp.conf \
-a -q -o \"%s:%R:%.2F\" |" );
while( <DIGITEMP> ) {
    chomp;
    if( $_ =~ /^nanosleep/i ) {
        my $now = localtime;
    } else {
        my ($dvc,$sn,$temp) = split(/:/);
        print "dvc=$dvc sn=$sn temp=$temp\n";
        $temp = 185 unless ($temp);
        for $i ( 0 .. $#sensor ) {
            if ($sn eq $sensor[$i]->{sn}) {
                $sensor[$i]->{temp} = $temp;
            }
        }
    }
}
close( DIGITEMP );
my $cmd = "rrdtool update temp.rrd N:";
for $i ( 0 .. $#sensor ) {
    $cmd .= ":$sensor[$i]->{temp}";
}
}
```

This can be set up as a cron job with a crontab entry:

```
*/* * * * * root test -c /opt/mh/local/bin/temp2rrd.pl \
&& /opt/mh/local/bin/temp2rrd.pl >/dev/null
```

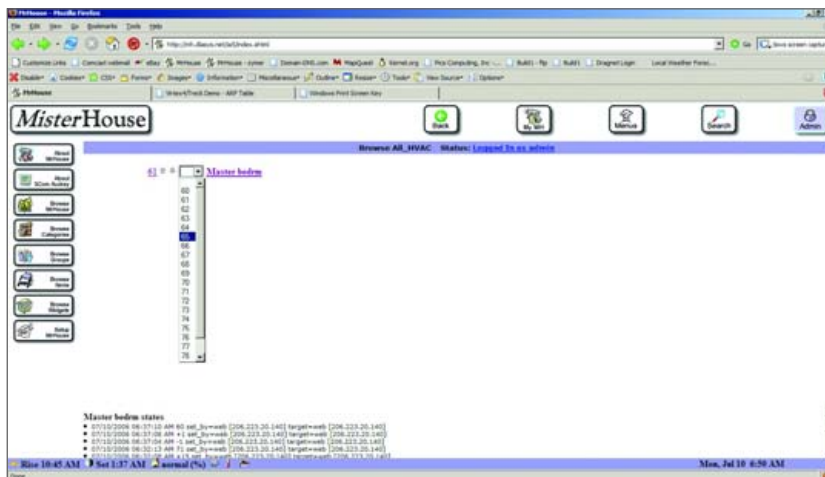


Figure 5. Browsing the All_HVAC group containing a single master bedroom zone—this shows a simple Web thermostat control to increase, decrease or set the setpoint for each HVAC_Zone object.

or it can be run periodically with MisterHouse.

I added a Perl routine to mh to query the RRD database periodically and compare current temperatures for each zone with the setpoint for that zone (specified in the mh UI code). If the current temperature is too far below the setpoint, the X10 module controlling the circulator for that zone is turned on. If the current temperature is too far above the setpoint, the X10 module is turned off. Right now, fault tolerance is achieved by turning on/off the appropriate X10 module on each loop if the current temperature is outside the acceptable differential from the setpoint. Eventually, the temperature sensors attached to the supply and return for each zone will be used to verify whether a zone circulator is running. They also will be used to calculate the actual heat output into each zone.

The RRD database keeps historical data, which might have a number of uses. Furthermore, RRD is designed specifically to record periodically polled values. It has configurable provisions to validate data and to handle queries against invalid data appropriately.

Perl is an object-oriented language, and MisterHouse takes advantage of that. Primarily, MisterHouse manages objects. The display elements, as well as programmed events and actions, all act on objects. MisterHouse already has a variety of interesting objects, but I needed something equivalent to a heating zone, so I created an HVAC_Zone object using an example object as a guide. My HVAC_Zones have a setpoint and a delta property. The setpoint is the target temperature for the zone, and the delta is the allowed deviation before taking action. HVAC_Zone objects inherit from the MisterHouse Generic_Item, so they appear in the UI under browse groups as well as in a few other places. Zones are created much like the X10 items described above. They can be created in the same .mht file with X10 items or in their own .mht file. Other HVAC groups might be North, Night and Vacation.

An example entry would be:

```
HVAC_ZONE, H1, master_bedrm, All_HVAC|Thermostats, 72, 2.
```

This would create an HVAC_Zone object named master_bedrm in groups All_HVAC and Thermostats, with a setpoint of 72 and a delta of +/- 2°. A few minor alterations to the MisterHouse Web server were necessary to create a Web control that is displayed for each zone. I used the X10 Lamp item as a reference, as it already had a drop-down as well as arrow buttons for relative changes.

Finally, I use a Perl script executed periodically by MisterHouse to query the current temperature of each DS1820 within each zone, comparing it to the current setpoint/delta for that zone turning on the appropriate X10 appliance module for that zone if the temperature within the zone is more than the delta below the setpoint. Currently, I repeatedly attempt to turn on the appropriate circulator so long as the temperature remains below the set range, and I repeatedly turn off the circulator so long as the temperature exceeds the set range. This provides limited fault tolerance. Turning an appliance module off when it may be on already accomplishes a small degree of fault tolerance. Before the next heating season, I expect to have code to verify the state of the circulator using the supply and return temperatures, as well as raise an alarm if the temperature in a zone drops more than twice the delta below the setpoint. ■

Resources for this article: www.linuxjournal.com/article/9193.

David Lynch is an underutilized software consultant engaged in sometimes fruitless attempts to make a living developing embedded and systems software—usually under Linux. In another life, he is an architect, and he currently keeps himself occupied when not randomly turning lights on and off throughout the house or writing software for clients by building his own home. Home automation combines his vocation of software with his interest in buildings, as well as satisfies his urges for world domination—at least within his own home. Dave can be reached at dhlui@dlasys.net and www.dlasys.net.

ApacheCon

October 9-13 • Austin Texas

US 2006

The Official Conference of The Apache Software Foundation

The premiere developer event showcasing leading Apache projects and Open Source innovation, ApacheCon is the ideal forum to meet and exchange ideas among peers, ASF members, vendors, and power users.

Experience the future of Open Source.

- Unparalleled educational, networking, and social events
- In-depth tutorials, presentations, and hands-on sessions
- "First-looks" of groundbreaking new technologies
- Internationally-recognized speakers and instructors
- Apache from A to X—key issues, tips and tricks, case studies, business track, and more.



Tutorials: October 9-10
Conference and Exhibits: October 11-13
Hilton Hotel • Austin, Texas


Register today!
us.apachecon.com

Do-It-Yourself Robots with Linux

Linux-based robots are tricky to create, but Michael Surran's Robotics class found out it can be done. **MICHAEL SURRAN**

Robots have been a passion of mine since I was a child, so imagine my excitement when I was given the opportunity to add a robotics class to our high school's computer curriculum! We recently celebrated our second year of offering robotics at Greater Houlton Christian Academy

(GHCA), the school where I teach. During this time, we've produced three different robots, each based on a PC running Linux. We work with a tight budget, so we have to be creative in our design, use of materials and tools. This results in robots that any do-it-yourself hobbyist can build.






Figure 1. K9 was the first robot built at GHCA.

The first robot to roam the halls of GHCA is K9, a robot dog based on the British sci-fi show, *Doctor Who*. I created K9 myself as a way to jump-start our Robotics program. K9's main purpose is to "ooh and ah" students, hopefully encouraging an interest in the field of robotics. I used common angle iron and 1/4 nuts and bolts to assemble K9's "skeleton", much like an Erector Set from years gone by. In fact, K9's head was built using an Erector Set from 1971! Sheet metal provides the finished look for our class mascot.

The second robot was designed by students Aaron Bither and Brian Thompson during our first year of offering robotics. At the beginning of each school year, we start by defining a "problem" that we'd like to overcome, and then we proceed to brainstorm how a machine can solve this problem, what features it will need and so on.

Aaron and Brian decided to build a robot that could deliver messages from one classroom to another. They quickly learned that what is simple for humans is usually quite complex for robots. By the end of the school year, they had a robot that could be programmed to navigate the hallway, but much was left undone. It was a learning experience for us all.

Our third robot, and the featured robot of this article, was designed by students Jordan McGuire, Jeromy Nevers and Barrett Jewell. I am also part of the team, contributing ideas and guiding the students, teaching them through the hands-on process of design and construction. The problem this group of students wanted to tackle is security, so they decided to make a "sentry bot" that could patrol the hallways and scan for intruders. Again, easier said than done.

A functioning, autonomous robot has three major components: the mechanical systems, the electronic systems and the software. Design typically begins with the mechanics of the robot, and of particular interest is the method of locomotion. There are many ways to make a robot move, but our class adheres to the KISS principle—keep it simple, silly! For example, we don't build complex legs when wheels work just fine. In fact, all three robots use ordinary lawn-mower wheels with various drive systems.

Our sentry bot uses a very simple tricycle design. To drive a robot's wheels, we use windshield-wiper motors. These motors can be purchased on-line for less than \$20 US or torn out of an old car for free. They run on 12VDC and provide very high torque at a nice rotation rate. Our sentry bot uses two of these motors: one to drive the front wheel and the other to turn the steering strut.

Once the means of locomotion has been chosen, it's time to design the




Figure 2. The Sentry Bot

frame. The frame needs to be strong but not too heavy, and it must accommodate the drive system, any additional mechanical systems (such as arms and sensors), the power source (we use a 12V garden tractor battery) and the electronics. We decided on plywood for the sentry bot because it is relatively inexpensive and easy to work with. We cut out the frame using a simple jigsaw. One trick to reduce weight was to cut out sections of the interior of the plywood frame—just don't cut out too much or the frame will become weak.

Our electronics systems consist of the main computer and interface board that connects the computer to the drive system and sensors. Although I would love to equip every robot with a low-power embedded computer, our budget simply does not allow for this. However, we do have

access to a number of old Pentium desktop computers. Considering the processing power of the Mars Rovers, a 200MHz Pentium is more than enough for our humble robot.

A big challenge in using old desktop computers for robotics is the power consumption. These motherboards draw a lot of current. Another problem is the voltages that the motherboard needs.

For the sake of time and money, we use a power inverter with the original AT power supply. This is easy to do but not very efficient. Future robots will use components from old laptops that are donated to us, which consume less power and can run directly off the battery.

Interface boards are designed and built in class. The main purpose of this component is to take the low-current outputs of the PC's parallel port and allow those outputs to power the motors. We use transistors and relays to construct our control circuits, as well as a logic circuit to prevent motor "twitching" during bootup. Our logic circuit also ensures that the motors don't exceed any limits, such as turning the steering strut too far left or right. Inputs from sensors are routed through buffers before going to the parallel port. Simple, but it works!

Of course, what you really want to learn about is the software and operating system that run our robots. Yes, all three robots run Linux! In the beginning, we experimented using DOS and QBasic. However, we were quite limited in what we could do with DOS, so I started experimenting with Linux using K9 the robot dog.

When K9 was built, we used the Mandrake Linux distribution at our school. Mandrake offered a minimum installation option, which was enough for the simple demands of our robots. Since then, however, our school has switched over to Gentoo Linux, and thus our sentry bot actually

Considering the processing power of the Mars Rovers, a 200MHz Pentium is more than enough for our humble robot.



Figure 3. Barrett Jewell works on the interface board.

runs on Gentoo. Maybe some of you are thinking that Gentoo is overkill for a simple robot. Maybe, maybe not.

I personally find that installing Gentoo is faster than installing most other distributions, maybe because I've done it so many times. A bare system (stage 3) install of Gentoo is relatively small, at least compared to the storage available on the robot's hard drive. However, Gentoo's biggest offering is the ability to add and update the software our robot needs quickly and easily, all compiled to make our old hardware run as fast as possible.

In case you are wondering how useful it is to compile software on an old Pentium system, know that our sentry bot has network access to our computer lab. Using distcc, a distributing compiler, our bot has the processing power of 20 Athlon 1500+ processors and more than 7GB of RAM at its disposal for compiling new software. Trust me, sentry bot can compile source code pretty quickly.

Gentoo Linux comes with another plus—it is very Python-centric (due to the Portage system). Our school offers a class to all 8th graders where they use Python to learn the basic concepts of programming. This provides a natural stepping stone for students who later want to take our Robotics class. Python is both simple and powerful, and although it is an interpreted language, it runs plenty fast for the simple needs of our robots. We use some C to interface Python to the parallel port, as well as to deal with time-critical operations, such as drift corrections. However, most of our code is in Python. This allows the students to make changes quickly and see immediate results in the robot's behavior. Because our computer lab is Linux-based, students easily can access the robot using SSH and KDE's fish from their workstations.

Controlling the motors using Linux and Python isn't difficult. The real challenge lies in the implementation of awareness, both of the environment and the robot's internal conditions. For example, a simple program coupled with a few electronic components can turn on the main drive motor for ten seconds, propelling the robot forward. But what's to stop the robot from drifting to the left or the right, or from running into a wall or person? Although the robot easily can track its progress based on time (driving ten seconds forward), the speed of the motor varies with battery charge and friction. To tell the robot to go forward exactly 50cm requires something a bit more complicated; it requires a feedback loop.

Most robots use motors with encoders of some sort to act as a digital "odometer". These encoders can be built using simple IR emitter/receiver pairs and a disk with holes in it. In fact, old-style computer mice work using this very principle. This got me thinking—instead of building a fancy encoder circuit and interface to the computer, followed by the software needed to read the encoder, why not use a simple device that provides the same information in a very computer-friendly format? Hence the

Expert Included.



Our experts provide dedicated customer service, backed by knowledge and ability based on years of experience. Art likes systems based on the Next-Generation AMD Opteron™ processor because he knows that its low power consumption and low heat output help keep costs down without sacrificing processing power or efficiency. AMD

Opteron processors offer high-bandwidth, energy-efficient DDR2 memory, and hardware-assisted AMD Virtualization™ streamlines deployment and migration, which Art knows is important to his customers.

When you partner with Silicon Mechanics you get more than great AMD performance — you get an expert like Art.



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



Figure 4. A Close-up of Our Optical Mouse Encoder

“invention” of the optical mouse encoder.

Like everything else about our robot, our optical mouse encoder is simple yet effective. We took a mouse pad, cut out a circle and glued it to the main drive wheel. We then mounted an optical mouse to the steering strut so that the mouse rests over the circular mouse pad. As the wheel turns, the mouse pad rotates beneath the mouse, which sends very precise positioning data to the computer. This movement easily can be read by opening and reading from the `/dev/psaux` or equivalent device file. Optical mice provide a high resolution with great accuracy, so with the proper conversions, we can track our robot's movement down to a fraction of a centimeter. Knowing how far it has traveled, the robot can know where it is on a preprogrammed map, assuming it hasn't been picked up and moved and that there is no drift or wheel slipping.

Of course, drift and wheel slipping do occur, and maps can't predict moved furniture, people or bookbags. The robot needs a way to detect its environment, and to do this we are using sonar and bump sensors. Typically, sonar is neither cheap nor simple, but it just so happens that I had recently “won” a car parking assistant at a Christmas party Yankee swap. This device uses sonar and has three lights to tell drivers how far to back up their cars as they pull into the garage. Well, you can't beat free, and although this device does not measure distance in the traditional sense, it was a working sonar that can detect objects that are close (yellow light) and really close (red light). By wiring the yellow and red LEDs to our parallel port's input lines, our robot can detect objects as it approaches them and take action, such as stop or turn. The sonar emitter is mounted on the front wheel strut, so the sonar is always “looking” in the direction the robot is heading.

Even though we've tried to keep things simple, we still have much work to do. Our students are looking forward to the coming school year when they can continue working with the physical platform they've constructed. With the majority of the hardware in place, our focus will turn to the software. One challenge we've already run into is making sure the robot is going perfectly straight and not drifting to the left or right. Although it may be possible to design the hardware to lock the front wheel pointing *exactly* forward, there is still the issue of drift due to friction differences at each wheel. To solve this problem, we're turning to software.

Instead of measuring the rotation rate of the single front wheel per our current design, we recently experimented by attaching two optical mice to measure the two rear wheels. Linux provides a separate device file for each mouse, allowing us to track both rear wheels independently and compare them. By measuring differences in rotation rate, we can use software to adjust the steering motor until the robot is going perfectly straight. This also will let us precisely calculate the robot's heading as it makes turns or is affected by drift and wheel slippage. Couple this with external sensors,




Figure 5. A Look at the Steering Mechanism and Interface Board

and we will have a robot that can know exactly where it is in our building.

High school student Jordan McGuire has written software that allows the robot to correct its internal position based on sonar readings of known boundaries, such as walls. Similar techniques can be used to detect objects that “don't belong”, and then use other sensors, such as heat and motion, to detect a possible intruder during times when the building should be empty. With Linux, it is simple to add a Webcam that can capture and send the images of an intruder via a wireless network. A sentry bot's best weapon is its camera.

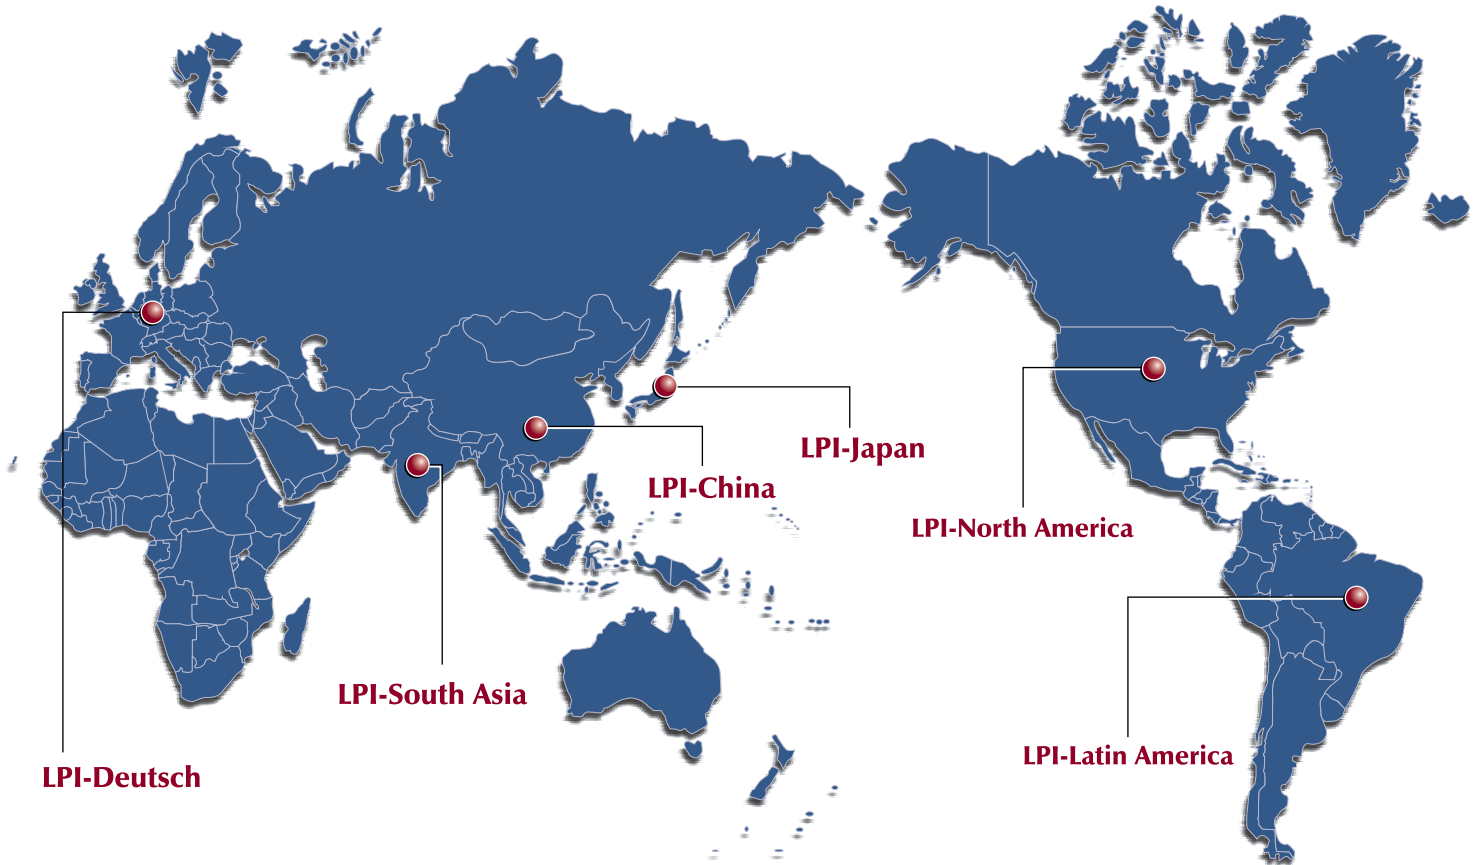
With Linux as our operating system, we can expand the capabilities of the robot well beyond the original “mission” given it. We've even discussed adding a GPS receiver and programming the robot to fetch us milkshakes from the local ice cream stand. Okay, maybe that's a little far-fetched, but it is good to dream big!

Speaking of dreams and future designs, as the Robotics program matures and students are introduced to robotics at a younger age, I plan to bring in more advanced concepts. We will still work to keep it simple, but there definitely are better ways to do some of the things we are currently doing. For example, the power requirements of a desktop PC are unrealistic for a battery-powered robot. Laptop components will be a great improvement, but if money were no object, embedded computers would be the way to go. A recent donation has purchased a TS-7260 ARM computer board for K9 that requires less than a watt of power to run. This Linux-based board has a number of I/O ports that are better suited for interfacing to the various motors, servos and sensors of a robot when compared to a PC's parallel port. The other inefficiency our simple approach introduces is the burden of a central PC managing minute yet very time-sensitive tasks, like speed control of the main drive motors. I'm currently redesigning K9's controller board to use simple microcontrollers to unload the low-level tasks from the CPU. Although our class may not be ready to tackle microcontrollers just yet, they definitely will be introduced to the concept and get to see the results.

I realize this article doesn't give step-by-step instructions on how to build a DIY robot. There are many good books and informative Web sites on that topic. My goal is to share what can be done with inexpensive, easy-to-obtain parts and our favorite operating system. Linux is loved by tinkerers and hobbyists, and an exciting and challenging extension to this line of tinkering is to add motors and sensors and autonomy. Like peanut butter and jelly, Linux and robotics are a perfect combination!■

Michael Surran has been the network administrator and computer science teacher at Greater Houlton Christian Academy for the last seven years. He has published two articles advocating the use of open-source software in K-12 schools. Outside his life of computers, he enjoys the great outdoors of Maine. Comments can be sent to michael@ghca.com.

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.



Stable. Innovative. Growing.

Building a Digital Lifestyle with Open-Source Technology

Mike Diehl gives an overview of the hardware and software he used to create a magnificent home digital lifestyle.

MIKE DIEHL

OVER THE YEARS, I've enjoyed the benefits of a lot of technology in my home. (Translation: nerds have all the cool toys!) Even though I have friends and family that don't have Internet access, I often marvel at the conveniences afforded by modern technology, much of it open source.

A few years ago, we ran network cable throughout the house, enabling us to access the intranet, as well as the Internet from any room in the house. Aside from being cooped up in the hot, dusty attic for hours at a time, it really wasn't as hard as you might think. Once we had ubiquitous network access, it only made sense to configure some services for use by the family. First, I set up Samba so that my wife and I could have a common place to store family documents, pictures and music. This way, my wife can use her Windows workstation to access our files, and I can use my Gentoo Linux workstation via NFS and "it just works".

Eventually, we set up an Apache Web server to serve up Web pages. I've got it configured to serve an intranet Web page, which is just a bunch of commonly accessed hyperlinks, and a different set of Web pages to visitors from the Internet, where I try to promote various side businesses.

Of course, I run my own e-mail server, exim for SMTP and Courier IMAP. I use dyndns to provide my server with a Fully Qualified Domain Name (FQDN) that is accessible from the Internet. This way, my e-mail and Web addresses don't change, even if I change Internet providers.

MythTV has got to be the neatest toy I've ever come across. MythTV is a digital video recorder, DVR, that runs under Linux. Basically, to build a MythTV box, you start with a standard Linux box, add a video capture/tuner card, a video card with S-Video or composite outputs and a whole lot of hard drive space. What you end up with is TiVo on steroids. With MythTV, members of the family can choose TV shows that interest them from a variety of different menus, schedule when and how they are to be recorded and watch what they want, when they want. The system can even record programs into particular groups. This way, my kids' shows get recorded into the "Kids" recording group, and my boys know how to use the remote control that came with the tuner card to navigate to their recording group to watch the shows that we recorded for them. The kids don't have to wait until Saturday morning to watch cartoons and educational programs, and we don't have to worry about them stumbling upon programs that aren't appropriate for them. MythTV downloads the local TV schedule over the Internet, automatically—all this with the added bonus that MythTV allows us to skip commercials. Not only can I skip commercials, but I also can increase the playback speed. By skipping commercials and setting the playback speed to 150%, I can watch a 30-minute newscast in 15 minutes or less.

But MythTV goes beyond simple DVR functionality—way beyond. We use the MythMusic module to play our CDs and

Basically, to build a MythTV box, you start with a standard Linux box, add a video capture/tuner card, a video card with S-Video or composite outputs and a whole lot of hard drive space. What you end up with is TiVo on steroids.

MP3s. All of our family pictures are on the computer, so we can view them on the TV using the MythImages module. My kids get a kick out of seeing themselves and people they know on TV. The MythVideo module allows us to watch .avi and .mpg video files that we download from the Internet. When you consider that MythTV includes a Weather, Games, News and DVD module, you start to see that MythTV is an all-around entertainment system.

Thanks to Vonage and their annoying commercials, almost everyone has heard of Voice over IP (VoIP). Most people know that you can use your computer and VoIP to talk to other people over the Internet for free. Some people know that you can use VoIP to talk to other people using regular telephones. But even fewer people know how easy this is to set up with an open-source software package called Asterisk. Asterisk is an all-inclusive telephony toolbox, which is just a fancy way of saying that anything you want to do with respect to telephones and phone calling, you can do with Asterisk.

At our house, Asterisk handles all incoming and outgoing calls and replaces the caller ID boxes and answering machines that we used to have scattered throughout the house. But our Asterisk installation goes beyond mere answering-machine functionality. It also screens our calls, preventing virtually any telemarketer from bothering us. When calls come in, the computer checks to see whether valid caller ID information is available. If it's not, the computer picks up the call and asks callers to enter their phone numbers. If they don't, the system hangs up on them. Once a valid caller ID has been received, the caller is presented with an answering machine greeting. The system is smart enough



Don't complicate a simple task

Keep basic tasks just that with handheld, stationary and vehicle-mounted wireless data collection from AML. While others are busy reinventing the wheel, we're keeping things simple, from our products to our personal service. Visit us at www.amltd.com

or call 1.800.648.4452 for a real live person.



M7100 Wireless Family

AML
NEVER COMPLICATED™



to greet our friends, family and coworkers by name. The greeting informs callers that they must press the # key in order to reach us; otherwise, they can press the * key or wait 15 seconds to be sent to voice mail. While this is going on, the computer displays the caller ID on the MythTV and announces the caller's name on the server's speakers. All this happens before our phones even begin to ring!

Since we installed the system, we've not had a single telemarketer bother to press the # key to talk to us. Also, many of the people who call us simply want to leave a message anyway, such as our dentist's office confirming an appointment.

We gained several features that would have cost us extra from our old phone company. Call waiting allows us to answer an incoming phone call even when we're already on the phone. The do-not-disturb feature allows us to tell the Asterisk system simply to send callers to voice mail, so we can enjoy some time at home without being interrupted by a ringing phone. Of course, we're able to access our voice mail remotely, either through the telephone network, or via e-mail. I even get paged when someone leaves us a new voice-mail message.

Because the people I work for are gracious enough to allow me to work from home on occasion, I tend to be on the phone quite a bit. The Asterisk system lets me have a separate virtual phone line in the office. Thus, I'm able to be on the phone without tying up the regular home phone. My wife can use the phone any time she needs to, even though I may be on the phone all day long taking care of business.

Asterisk and VoIP give us a lot of flexibility as to how we use our phone system. We can use a VoIP phone such as a Cisco 7960, or we can use one of the many soft phones like X-Ten. Because Asterisk supports all of the major VoIP protocols such as SIP, H.323, MGCP, Skinny and IAX2, our options are virtually limitless. We also have the option of installing an Analog Telephone Adapter (ATA), which allows us to use our existing telephones and telephone wiring. But one of the best features of our VoIP system is that it's a lot cheaper than our old PSTN phone line. When we were with Qwest, we had an unlisted phone number and caller ID, which cost us \$40 US each month including long distance. Because I'm able to buy VoIP phone service wholesale and add call features such as voice mail, our last phone bill was about \$15 US for the same or more phone usage.

Of course, all of these services place a high demand on disk storage. It's not enough simply to buy a bunch of hard drives. You have to have some mechanism of organizing them in a meaningful fashion. Because I have more than one terabyte of disk storage, I use the Enterprise Volume Management System (EVMS) to manage my hard drive space. Though I could have opted for any conceivable RAID configuration, I opted for a simple linear drive append configuration. This essentially amounts to being able to add up the storage of all of the installed 200–320GB drives and present them as one large drive on which I can put one or more filesystems. More important, EVMS gives me the flexibility to move data from one physical disk to another without having to move it from one filesystem to another. I essentially have four physical hard drives that combine logically to create one very large, virtual hard drive. Drives /dev/hdb1, /dev/hdc1, /dev/hdd1 and /dev/hda4 combine to form /dev/evms/media, which contains a single 750GB filesystem.

This feature came to my rescue recently. A couple weeks ago, I started seeing hard drive I/O errors from one of the "middle" drives in my system log. In most cases, this would be bad news. But because I caught the problem early and was running EVMS, I was able to buy a new hard drive and get it installed in the server. Then, I was able to use the EVMS management utility to migrate the data off of the failing drive and onto the new drive, without having to change the overlying filesystem. I simply unmounted the filesystem, added the new drive to the logical volume and migrated the data off the failing hard drive. When it was finished, I simply remounted the filesystem with little or no loss of data. So by using the same principle, I plan on being able to retire older drives and replace them with newer drives, without having to mess with the usual backup and

restore cycle. The volume manager simply will move the data from the old drive to the new drive transparently.

But what happens when the day comes when I want to replace an aging 200GB drive with a new 400GB drive, but I don't have any empty drive slots? Obviously, I'm not going to be able to do a backup and restore on 750GB of data! Well, this won't be a problem because of a native Linux protocol called ATA over Ethernet, or ATAoE. Using EVMS, ATAoE and an open-source package called Vblade, I'll be able to install a new drive in my desktop computer and access it natively on my server computer. EVMS will see the "remote" hard drive as a locally installed drive and allow me to migrate data to it over the network. Once the migration is complete, I simply will remove the old drive and install the new drive into the server. When I reboot the server, it will be as though the new drive had always been there.

As many, if not most, of my friends also have home-based networks, it seemed like it might be fun to try to connect them all together. Because most of my friends have dynamic IP addresses and some of them use systems other than Linux as their Internet routers, IPv6 was out of the question. Instead, I use the open-source package, OpenVPN, to connect my friends' networks. Setting up a VPN connection between my friends' networks gives us all transparent access to each other's computers, even computers behind NAT firewalls. Ubiquitous access to each other's networks and computers allows us to do private file and resource sharing.

Using a program called Unison, we're able to keep a synchronized archive of family pictures as well as encrypted backups of each other's financial documents. Because of the dynamic nature of such a network of networks, we use the open-source package Quagga, which uses OSPF and RIP to manage the routes between the various networks. So far, this effort has been more of a learning experience than a practical benefit, but we're still working on it.

After spending so much time getting these toys and gadgets working, it's time to start working on some practical tools. For one thing, we're drowning in spam. I've got the Exim mail server configured to drop incoming e-mail into folders based on address and subject. This has worked pretty well at reducing the amount of time I spend reading e-mail. I can delete entire folders that don't interest me. However, it's clear that I need to do more. I'm considering configuring SpamAssassin to triage incoming e-mail messages. SpamAssassin performs a series of tests on a given e-mail message to determine how likely it is to be spam. Each message is given a score. The higher the score, the more likely the message is spam. Those messages that SpamAssassin deems to be probable spam will then be routed to TMDA for ultimate verification.

TMDA implements what's known as a challenge/response mechanism for authenticating e-mail. Essentially, for any message that TMDA handles, TMDA sends a reply

message to the original sender. This is the challenge part of what TMDA does. If senders are non-spammers, it's likely they will reply to the challenge message. This is the response part. When TMDA receives a response to a given challenge, it arranges for the original e-mail message to be delivered. The theory is that most spammers won't reply to the challenge, so TMDA won't deliver the spam message. Eventually, the spam will be expired and deleted, having never been delivered. The only problem with TMDA as the sole spam filtering mechanism is that some people don't want to be bothered by challenge messages. That is why I intend to run the messages through SpamAssassin first. This way, only spammy messages will be challenged.

I've tried to outline in this article some of the really neat technology available in the open-source arena. Hopefully, you've read about something that you didn't know about before. Now, I'm off to try to figure out why my VCR keeps flashing 12:00. Oh well. ■

Resources for this article: www.linuxjournal.com/article/9192.

Mike Diehl works for SAIC at Sandia National Laboratories in Albuquerque, New Mexico, where he writes network management software. Mike lives with his wife and two small boys and can be reached via e-mail at mdiehl@diehlnet.com.

Because I have more than one terabyte of disk storage, I use the Enterprise Volume Management System (EVMS) to manage my hard drive space.

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

Manage Your Photos with F-Spot

F-Spot hits the ease-spot to provide photo management for all.



DANIEL BARTHOLOMEW

BROWSE MODE TOUR

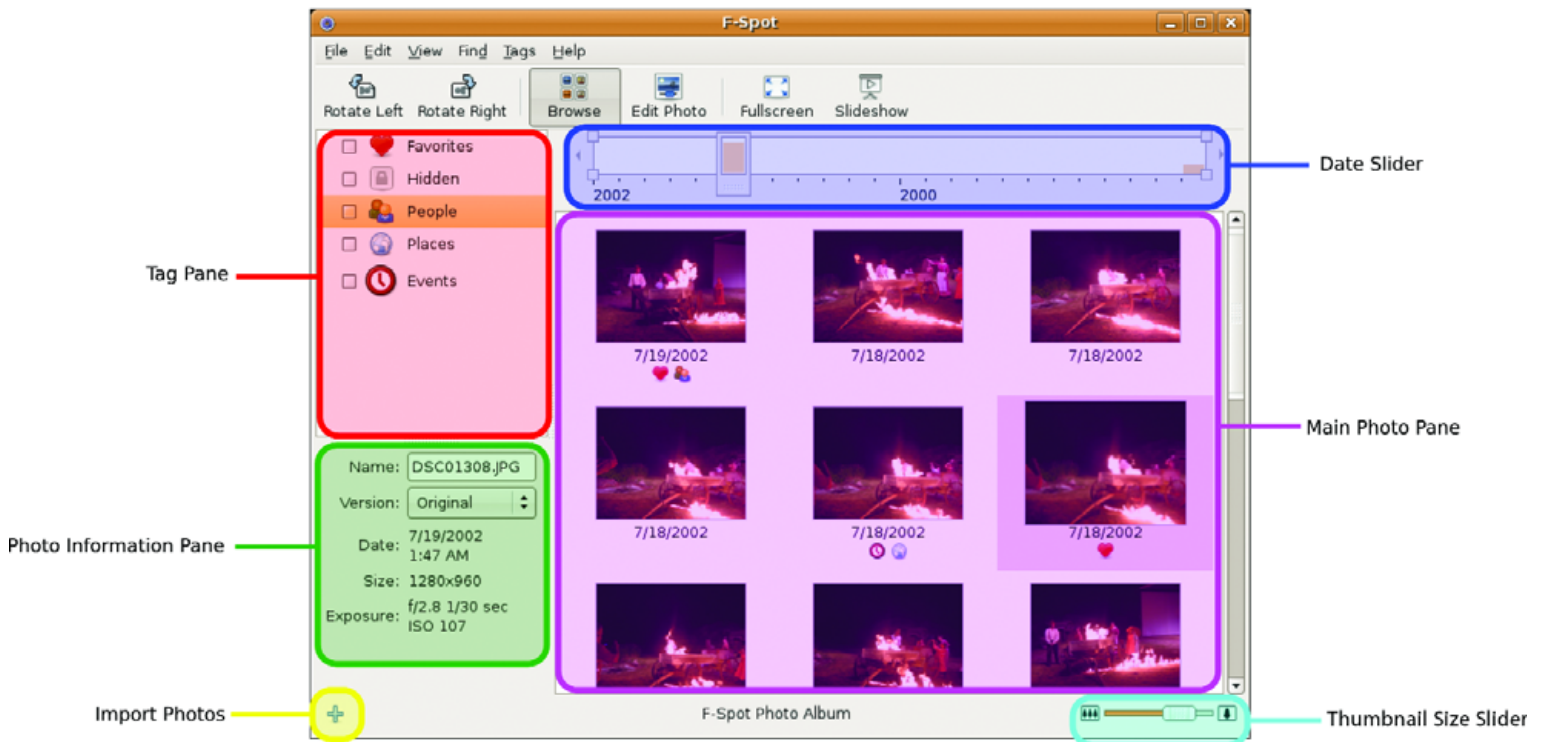


Figure 1. The F-Spot Interface

During the past several years, my wife and I have been trying to organize our ever-increasing collection of digital photographs. First, lacking a digital camera, all we had to do was organize the photos we had scanned. Because scanning was a tedious process, we scanned only the photos we felt were the most important—births, weddings and other life events. This kept the numbers low and manageable. However, once we purchased our first digital camera, our collection of digital photographs quickly grew out of control. With an effectively infinite supply of “film” we suddenly found ourselves “needing” to record each and every moment of our children’s lives.

For a long time, the organizational method (if you could call it that) was to take photos until the memory card on the camera was full and then pull them off of the camera and shove them into a folder with a name corresponding to the current date. This was less than ideal, especially because with a large memory card, we could take several hundred photos during a period of a few weeks before it filled up and became necessary to transfer them.

We also had no way to find, say, all of the photos of a particular child easily. Other than knowing when the photos were copied off of the camera, the photo collection was an unorganized mess. Thankfully, F-Spot came to our rescue in a big way.

F-Spot is a photo management application for the GNOME desktop. It is one of a growing number of Linux applications built upon Mono, the open-source implementation of Microsoft’s .NET platform.

The F-Spot program is available in the package repositories of most Linux distributions, including Fedora, Ubuntu, SUSE, Debian and others—basically any distribution that comes with a recent version of GNOME. Therefore, as a first step, check your distribution’s package repositories. Or, if you’re a command-line person, you can install F-Spot as easily as typing `apt-get install f-spot` for Debian and

its derivations, such as Ubuntu.

If installing from source is your only option, you can download the source from the F-Spot Web site (see the on-line Resources). F-Spot requires that you have the GNOME development libraries, Mono, SQLite, libgphoto2 and a few other libraries installed and configured on your system as specified in the README in the tar.bz2 source file. Once the prerequisites are taken care of, installing it is a simple `./configure; make; make install`; three-step process.

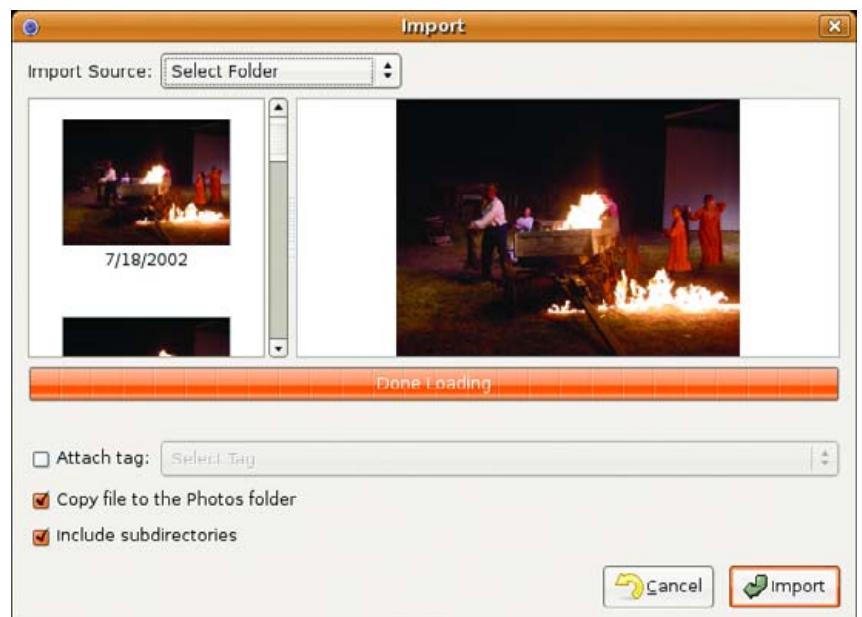


Figure 2. Importing Photos into F-Spot

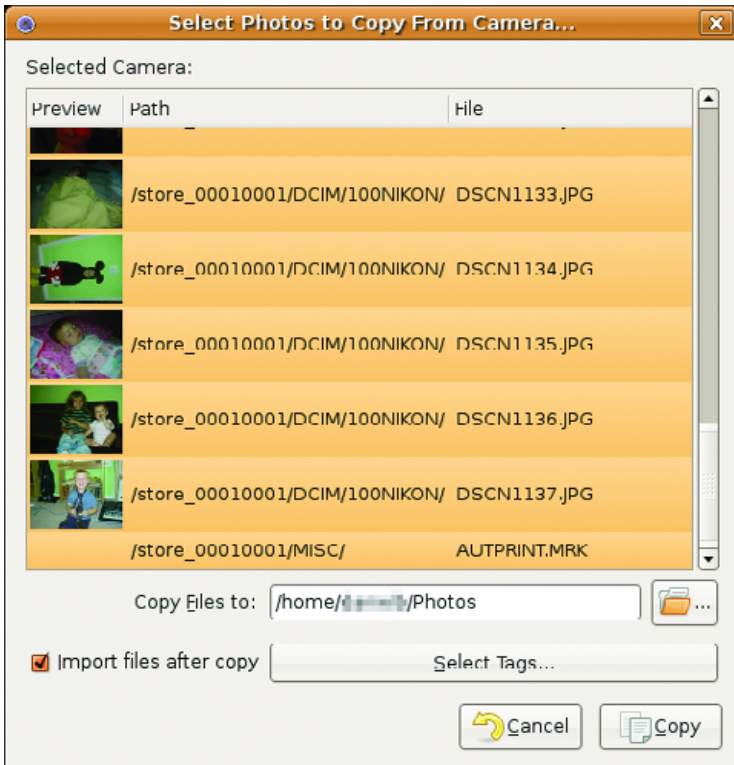


Figure 3. Importing Photos from My Digital Camera



Figure 4. Copying Photos from My Digital Camera

After installing F-Spot, the first task is importing existing photos. Choose import from the File menu and then Select Folder from the Import Source drop-down menu. This brings up a file selection dialog. Navigate to the folder where you have been keeping all of your photos, and click Open.

After you choose a folder from which to import, F-Spot searches through the folder and loads thumbnails of the photos it finds. Then, select the Copy file to the Photos folder and Include subdirectories check boxes, and click Import. In my case, at this point I sat back and watched as F-Spot imported and organized several years' worth of photos.

Importing from a digital camera is likewise a very easy process. Simply connect the camera to your computer, and choose the camera from the Import Source drop-down menu. F-Spot recognizes most digital cameras automatically.

F-Spot organizes your pictures in a folder in your home directory called Photos. It subdivides that folder into a folder for each year, and those folders have subfolders for each month and then another for

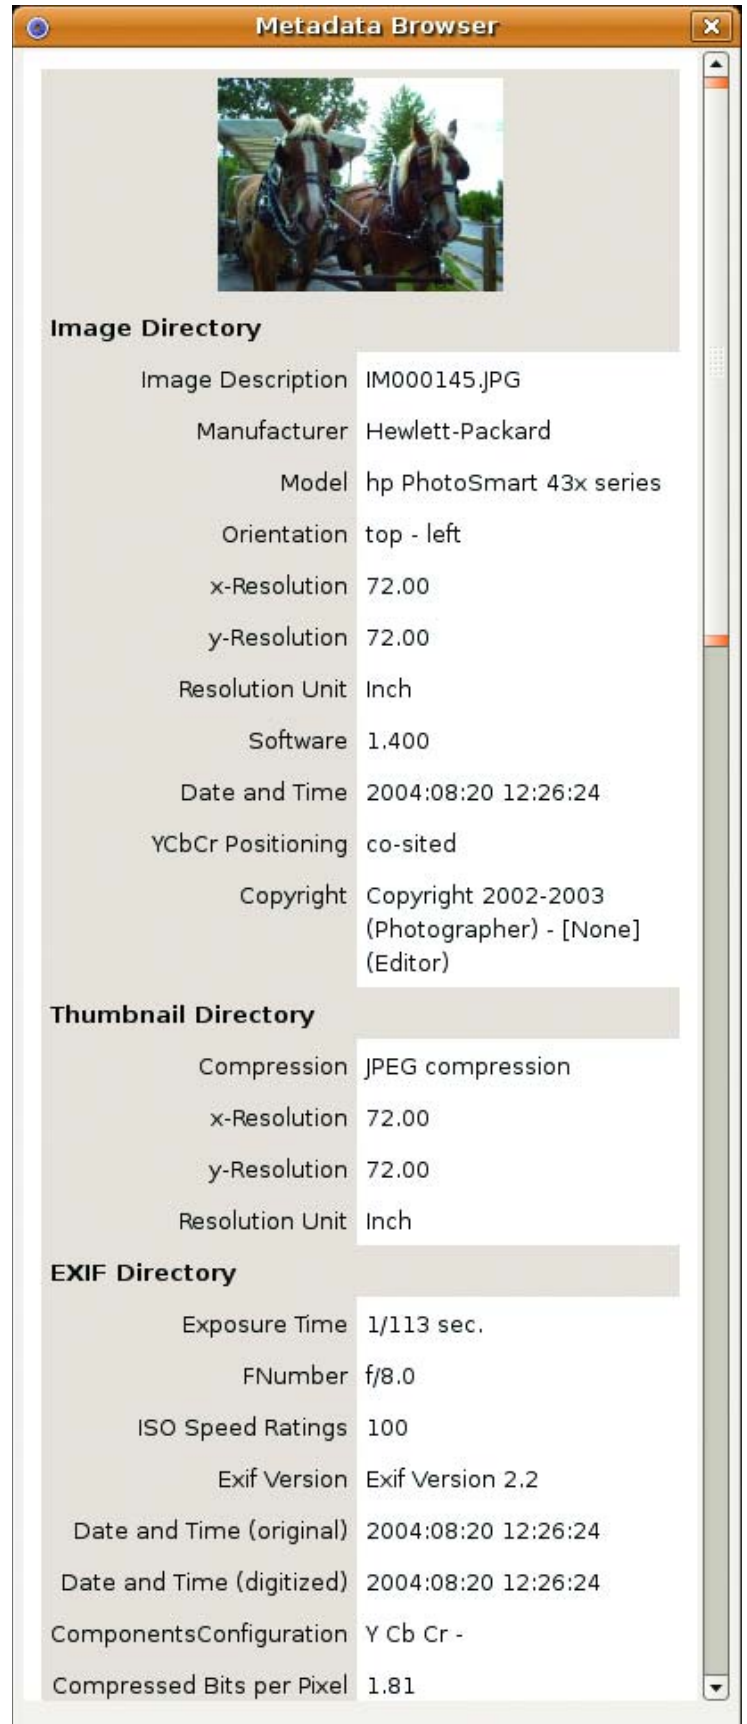


Figure 5. The F-Spot Metadata Browser

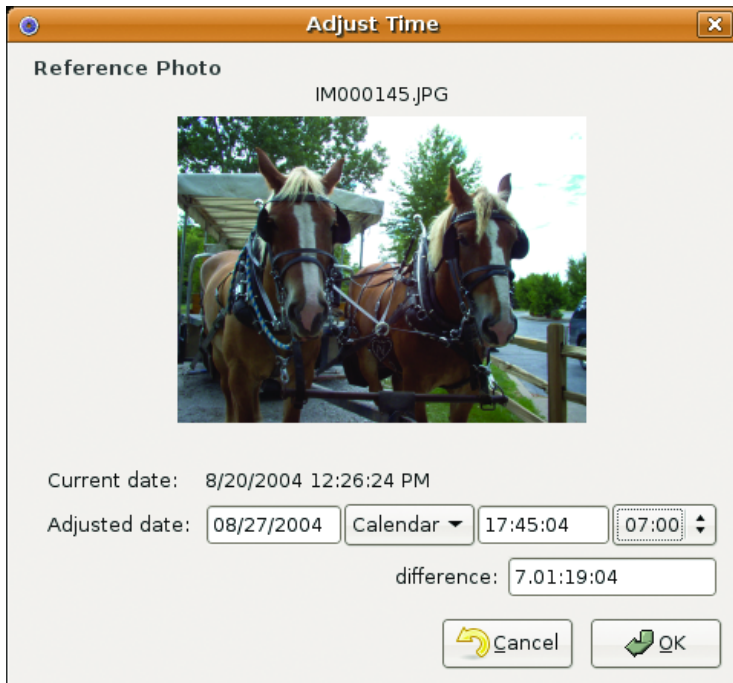


Figure 6. F-Spot makes it easy to adjust the time.

each day. The program figures out the appropriate dates by reading the EXIF metadata from the photos themselves. As long as the clock on your digital camera is set correctly when you take photos, F-Spot files it correctly.

Digital cameras record all sorts of information in the picture files they create, such as the f-stop used, whether the flash fired, the manufacturer and camera model, the exposure time and much more. You can use the Metadata Browser, found in the View menu, to look at this information. It is not editable, but it is interesting to look at, especially if you are a photography enthusiast.

If the metadata is missing for whatever reason, F-Spot sets the date using the creation date stamp on the file, which is usually not accurate. F-Spot can add the date metadata to the picture files that don't have it. This came in handy for us especially with the the 35mm photos that we have scanned over the years, which did not have any metadata associated with them. I also found several photos that, although most of the metadata was intact, had incorrect dates associated with them, probably due to an incorrectly set clock on the digital camera that took the picture.

Fixing the date is easy; simply choose Adjust Time from the Edit menu. This brings up a dialog for adjusting the date down to the second (if that's what you want). You also can select a group of photos, say, from a party, and have F-Spot space them five or ten minutes apart automatically (or however many minutes apart you choose), which ensures that they stay in the proper sequence.

Once you organize your photos more or less by date, you may want to tag them. Tags in F-Spot work like an extra layer of organization on top of the physical time-based scheme. F-Spot comes with a fairly generic set of tags: Favorites, Hidden, People, Places and Events. I added tags for each of my children, and also tags for clipart, landscape photos, desktop wallpapers and so on.

Tags can be nested as well. For example, all of the tags for my children are nested under a Family tag, and my Wallpapers tag is subdivided into Abstract, Landscape and Other.

Photos also can have more than one tag associated with them, so for photos with more than one child in them, I simply tagged them with every child in the photo. You can add tags by dragging the

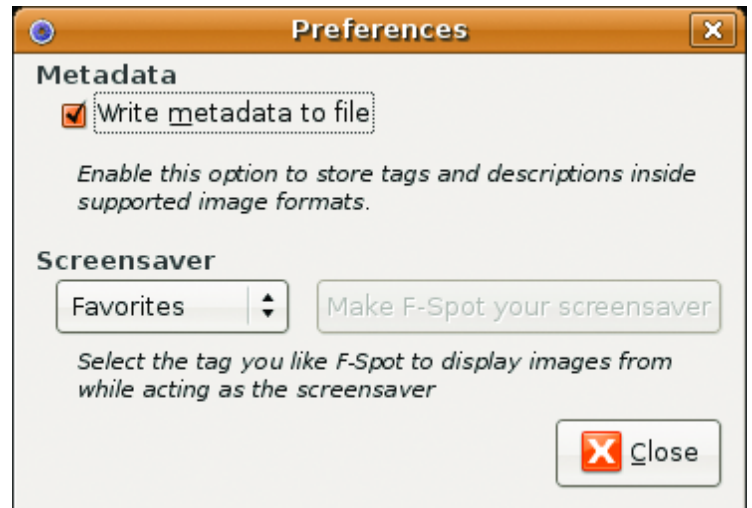


Figure 7. The F-Spot Preferences Dialog


appropriate tag to the photo or by dragging the photo and dropping it on the appropriate tag. To save time, you also can do this with several photos selected.


After tagging all of my photos, I can easily view only the photos I want, such as landscape desktop wallpapers, by selecting the check box next to the Landscape tag. If instead I want to look at all of my wallpapers, I simply check the parent Wallpapers tag.

Linux Laptops

Starting at \$990

**DON'T BE SQUARE!
GET CUBED!**





309.34.CUBED

shoprcubed.com

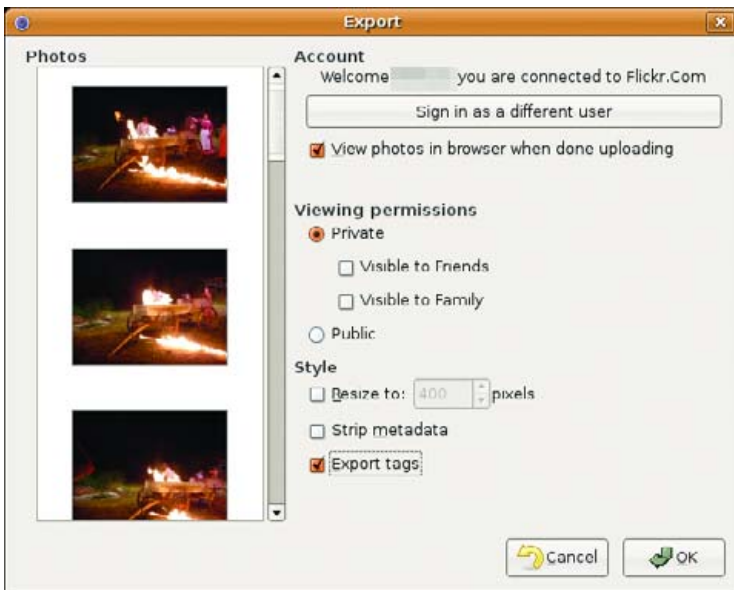


Figure 8. Exporting Photos to Flickr

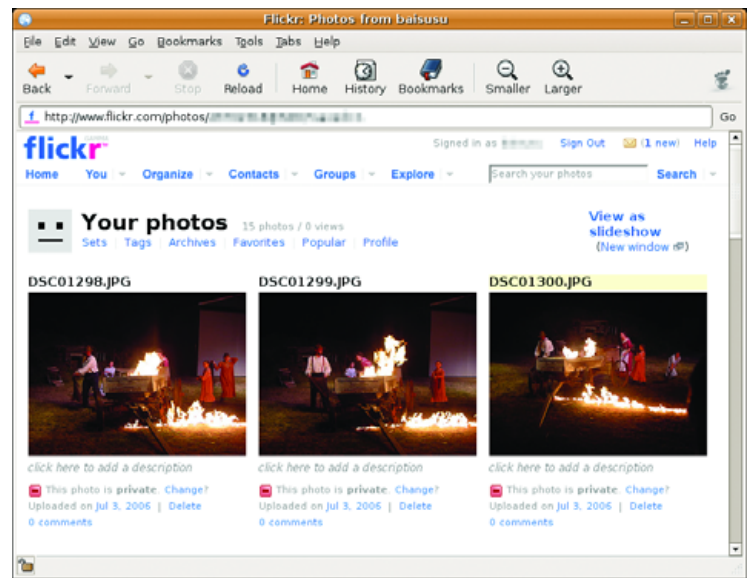


Figure 10. My Photos Uploaded to Flickr

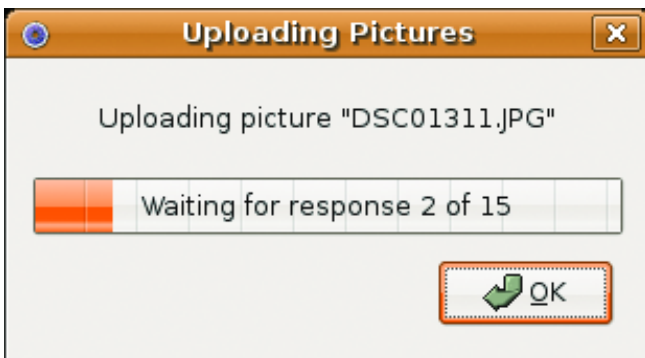


Figure 9. Uploading to Flickr can take a while.

The tagging feature also comes in handy with another of F-Spot's useful tricks. It can function as your desktop screensaver. Enable this mode via the F-Spot preferences dialog under the Edit menu. From the preferences window, choose which tags to pull photos from. The screensaver itself is nothing fancy; it is just F-Spot's Slideshow mode, which crossfades between the photos—no panning or zooming here—but at least, thanks to specifying the tag, you are guaranteed F-Spot displays only the photos you've chosen.

The preferences window has only one other function: to enable or disable the writing of metadata back to the image formats that support metadata.

In addition to serving as a screensaver, F-Spot has other neat GNOME desktop integration tricks. For example, to select a photo (or photos) to e-mail to someone, choose Send Email from the File menu, and F-Spot opens a new Evolution compose window with the images already attached. You then can compose your message and send it without having to launch the complete Evolution application.

Apart from desktop and e-mail integration, F-Spot also can interface with the Flickr Web site. Flickr is a photo storage and sharing Web site owned by Yahoo. Access this option from the Export submenu of the File menu. F-Spot takes you through a brief setup process the first time you use it and from then on it remembers you.

Uploading a large amount of photos to Flickr can take a long time, so

it is best to upload in small batches. Slowness aside, uploading to Flickr is fairly painless and much easier than uploading the photos manually.

You also can export your photos to a folder on your computer—as a basic but attractive photo album Web site, or simply as a folder of photos.

If you are looking for a simple way to back up your photos, you can export directly to a CD. The CD-burning feature uses GNOME's built-in CD burning facilities—another GNOME integration trick.

Like most photo management applications, F-Spot includes several basic photo editing controls. Enter editing mode by double-clicking the photo you want to edit or by selecting it and then clicking on the Edit Photo button on the tool bar. The editing features are basic, but very serviceable. Photo rotation, red-eye removal and cropping to various standard or nonstandard photo sizes are all easy to do, as is converting photos to black and white or sepia tone. The color adjustment window lets you tweak the exposure, saturation, brightness, hue, contrast, temperature and tint of your photos to your heart's content.

If you find yourself needing more editing power, right-click on an image and open it in The GIMP. F-Spot will even suggest that it create a new version of the photo to protect the original version.

In fact, if you change your mind about any of the edits you have made, you easily can revert back to the original version of the photo that existed when it was first imported into F-Spot. F-Spot does this by implementing a simple versioning scheme to protect the original and subsequent copies of the picture. Any time you make an edit, it does so on a copy, and you can create a new version of a photo manually at any time to tinker with it and not worry about irretrievably damaging the original or any prior version.

For all of its good points, F-Spot is not perfect. At the time of this writing, it is at version 0.1.11, and the application feels less than polished in a few areas.

Tag searching is one such area. As explained earlier, you search by tags by selecting the appropriate check box. It would make sense that by doing this on several tags, you could narrow your search to only those photos that are associated with *all* of the tags—sort of like what Google and other search engines do with keywords when I am looking for something on-line. Instead, F-Spot broadens your search to include all photos that have *any* of the tags. This can be useful in certain situations, but without a way to alternate between searching on all or any

EDITING MODE TOUR

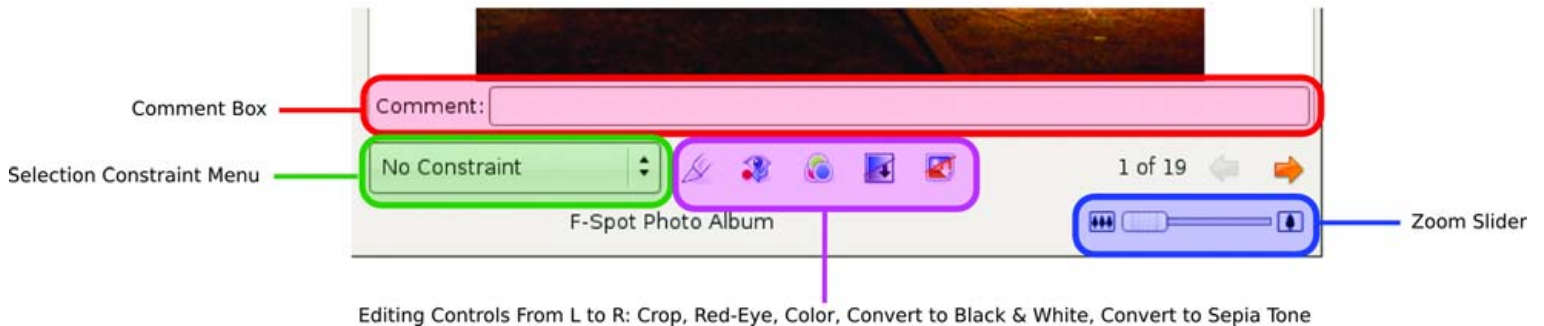


Figure 11. The F-Spot Editing Controls

of the selected tags, the search feature is, in my mind, broken.

Another annoyance is F-Spot's Gallery integration—it does not work, at least, not with the latest versions of each as I write this. Gallery is an open-source set of PHP scripts that you can use to add a photo gallery to your Web site. The integration with Flickr is nice, and it works well, but I'm not a Flickr user. I do like and use Gallery though, so the broken Gallery support in F-Spot was a let down. Hopefully, this is a temporary problem that will be fixed by the time you read this.

CD burning is another area that could be improved. It works, and works well, but not if you use 650MB or 700MB CD-R or CD-RW discs or, apparently, recordable DVDs. If you have only a couple hundred megabytes of photos to back up or 3GB, the disc images F-Spot creates are sized for 800MB CD-R and CD-RW discs only. There should be a way to choose what size media you intend to use and have F-Spot create its ISO images accordingly.

Finally, it would be nice if some of the photo editing features were a little easier to use. I can tweak the exposure, saturation, brightness, hue, contrast, temperature and tint of my photos, but I usually end up making my photos look worse, not better. Some automated "enhance photo" scripts would be nice or even something as simple as an "auto-levels" feature, such as you can find in The GIMP.

Despite these annoyances, F-Spot is off to a great start. It has not been around very long, and it already is one of the better Linux photo management programs. The GNOME integration and list of features already put it head and shoulders above gThumb, my previous choice for photo management in GNOME. As more developers add their support to Mono and its growing list of applications, F-Spot can only get better and better.

In the end, F-Spot became my photo manager of choice because it has brought order and sanity to my ever-growing collection of family snapshots. If you're a GNOME user, or curious about Mono, you owe it to yourself to give F-Spot a try. ■

Resources for this article: www.linuxjournal.com/article/9191.

Daniel Bartholomew has been using computers since the early 1980s when his parents purchased an Apple IIe. After stints on Mac and Windows machines, he discovered Linux in 1996 and has been using various distributions ever since. He lives with his wife and children in North Carolina.

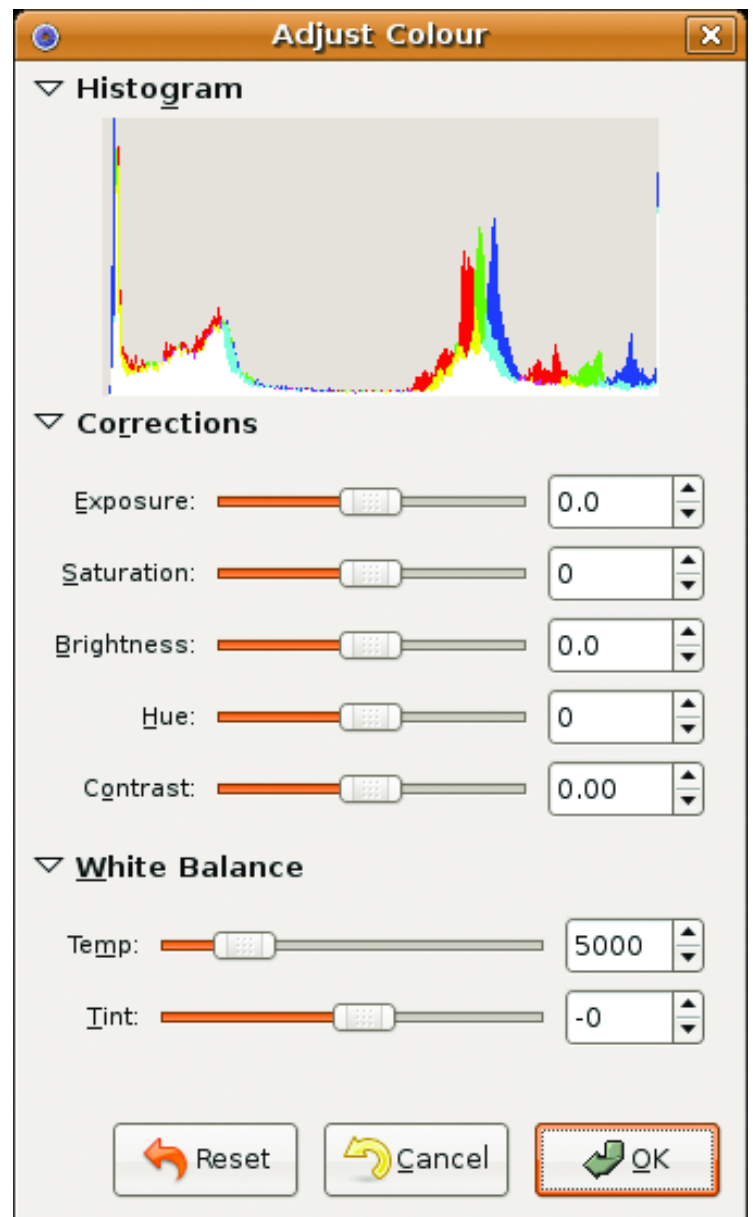


Figure 12. The F-Spot Color Adjustment Window

Advanced Video Coding on Linux

Use H.264 to create high-quality, low-bitrate digital video with currently available tools on Linux.

DAVE BERTON

The impact of H.264 on the world of digital video compression is growing. Companies such as Apple are already switching wholeheartedly to it. As part of the MPEG-4 standard (part 10), H.264 is now a part of both the HD-DVD and Blu-ray specifications for high-definition DVD. And for good reason—H.264 can encode video using very low bitrates while maintaining an incredibly high perceived quality.

Of particular interest are the low-bitrate possibilities this video codec provides. Luckily for those who run Linux, the H.264 codec (also known as the Advanced Video Codec, or AVC) has a successful and effective open-source implementation known as x264. In fact, the x264 Project won the Doom9 2005 codec comparison test (see the on-line Resources). x264 continues to make progress and improvements, and it remains an active project. So let's take advantage of what it offers us: an extremely high-quality AVC encoding tool that can be used right away for DVD and home movie backups, to create video clips for streaming over the Web or simply for experimenting with the latest video encoding technology.

The balance of this article focuses on the basic steps involved in creating standard .mp4 files that contain H.264 video coupled with AAC audio (Advanced Audio Codec, also an MPEG standard). The vagaries and subtle corners of hard-core video encoding are beyond the scope of this discussion. But hopefully, this introduction will encourage you to explore the topic further.

Because both AVC and AAC are now MPEG standards, it stands to reason that many tools (commercial and otherwise) are already available that support it. For example, Apple's QuickTime natively supports the video files we will be creating. And, MPlayer, the well-known and successful open-source media player, also supports .mp4 playback.

Getting Started

Creating a standards-compliant video file involves three basic steps: the creation of the encoded video, the creation of the encoded audio and the combination of those two things. Here are the software tools we need:

- MPlayer (includes mencoder, cvs version 060109 or higher)
- faac 1.24 or higher
- MP4Box (part of gpac 0.4.0 or higher)
- x264 (compiled with gpac support)

Our goal is to produce a low-bitrate video file suitable for posting on the Web. It will be a small file, but the quality will be exceptional compared with a higher-bitrate XviD encoding. Our source video will be a home movie clip called max.dv, which is a nine-second raw DV file captured directly from a digital video camera.

Let's process the audio first, as it is a pretty straightforward operation. The idea is first to have MPlayer dump the raw pcm

audio directly from our video source:

```
mplayer -ao pcm -vc null -vo null max.dv
```

This produces a file called audiodump.wav. The video portion of the source file is ignored. Now, encode this wave file to AAC:

```
faac --mpeg-vers 4 audiodump.wav
```

The --mpeg-vers parameter specifies the MPEG version. We now have the audio portion of our work finished and can listen to audiodump.aac by playing it with MPlayer.

When it comes to encoding the video, we are faced with several options. The highest quality encodes can be made only by using multiple passes. We actually process the source video twice (or more) in order to allow the encoder to pick the best possible distribution of bits across the destination file. Using multiple passes also enables us to pinpoint the bitrate and resulting file size of the output. However, encoding with an AVC encoder, such as x264, is very processor-intensive and thus can run pretty slowly, so we may not want to sit through a lengthy multipass encoding. Instead, we could run the encoding with one pass. This still will produce outstanding results, but never as good as a multipass encode. We also give up the possibility of targeting the resulting file size and bitrate. It all depends on what is most important to you, time or quality.

Fortunately, x264 provides a good middle ground. An option exists to specify a Constant Rate Factor (or Constant Quality), which instructs x264 to take into account the differences between high- and low-motion scenes. Because your eye loses details in high-motion scenes anyway, x264 uses fewer bits in those spots so that it can allocate them elsewhere, resulting in a much improved overall visual quality. This mode allows the highest quality possible without using multiple passes, which is a great time saver. The cost in using this mode, however, is in giving up the ability to determine the final file size and bitrate. Although this is possible with multiple passes, we would be forced to double the encoding time. So for our example, let's stick with one pass, utilizing the Constant Rate Factor feature (--crf) for greatly improved quality. Good values of the Constant Rate Factor range between approximately 18 and 26 (where a lower value produces higher quality but larger file sizes). Your needs in terms of size vs. time vs. quality may be different, however. If so, you should investigate multipass mode further to gain more control.

The x264 encoder accepts only raw YUV 4:2:0 input. To do this, simply pipe the output of mencoder directly into x264:

```
mkfifo tmp.fifo.yuv
mencoder -vf format=i420 -nosound -ovc raw -of rawvideo \
  -ofps 23.976 -o tmp.fifo.yuv max.dv 2>&1 > /dev/null &
x264 -o max-video.mp4 --fps 23.976 --crf 26 --progress \
  tmp.fifo.yuv 720x480
rm tmp.fifo.yuv
```

As you can see, we must specify the framerate (--fps); otherwise x264 will not know what is being fed into it. Do this similarly for the width and height of the incoming raw video. Encoding in this way enables the x264 default encoding parameters, which are quite good, but we can make a few improvements. In particular, we can make general improvements to some of the encoding strategies it uses without sacrificing too much in the way of extra encoding time. The number and variability of the parameters you can feed into x264 is great, and they are all geared toward improving the quality of the resulting output in some way. However, some options are more expensive, time- and processor-wise, than others. And, some options can sacrifice compatibility with certain media players, notably QuickTime. In order to remain compatible with the existing install base of QuickTime users, we need to keep a few things in mind.

QuickTime and H.264

It is nice that QuickTime 7 supports H.264-encoded video. Apple itself encodes all of its movie trailers on-line using H.264. Although this is good, and fosters the adoption of this codec, the QuickTime implementation has some limitations, most notably with B-Frames and Profile support. We need a short detour to explain what this means for our encoding project.

The MPEG standard for H.264 includes a number of profiles, such as Baseline, Main, Extended and High. These profiles delineate different technical capabilities that a decoder may need to possess. As its name suggests, the Baseline profile is the simplest and least-demanding profile, and Main, Extended and High require more processing power and the interpretation of more technical features in order to decode properly. QuickTime 7 supports Baseline and parts of the Main profiles; however, it chokes on features of the Extended and High profiles.

B-Frames are a type of storage format for digital video. These types of frames reference information from other previously decoded frames in order for the decoder to do its job properly, which is to decode the video. B-Frames are interleaved amongst other frame types known as I-Frames and P-Frames. It's a technical detail, but the QuickTime 7 H.264 decoder can support only up to two B-Frames, no more. This is unfortunate, because using more B-Frames would let us increase quality under some circumstances.

To remain QuickTime-compatible, we need to keep these limitations in mind. However, the quality of our low-bitrate encoding will not really suffer that much, even with these limitations. And, we can enable a few additional options to improve things quite a bit. The first is the subpixel motion estimation (--subme) size, which controls the precision of motion estimation calculations used by x264 during the encoding process. By increasing this to 6, the maximum, we gain a lot of visual quality at the cost of some additional encoding time, but it is worth it. We also can configure how x264 analyzes frames to perform better motion estimation (--analyse), which leads to higher-quality encodes. Note that some types of analysis are for High profile encodings only, such as 8x8 DCT, which are not supported by QuickTime, so we avoid those settings. We also can disable PSNR calculations (--no-psnr) to buy back a little speed during the encode. PSNR is simply a quality measurement and has no effect on the actual encoding quality.

Putting all this together, we can now output a high-quality, low-bitrate, QuickTime-compatible and standards-compliant video encoding using H.264:

```
mkfifo tmp.fifo.yuv
mencoder -vf format=i420 -nosound -ovc raw -of \
  rawvideo -ofps 23.976 -o tmp.fifo.yuv \
  max.dv 2>&1 > /dev/null &
x264 -o max-video.mp4 --fps 23.976 --bframes 2 \
  --progress --crf 26 --subme 6 --analyse \
```

```
p8x8.b8x8,i4x4.p4x4 --no-psnr tmp.fifo.yuv 720x480
rm tmp.fifo.yuv
```

We can make further improvements. Because this video file is destined for the Web, we most likely want to reduce the frame size to something more friendly, possibly crop out unwanted areas, and make other adjustments. For example, to reduce the frame size, run the following commands:

```
mkfifo tmp.fifo.yuv
mencoder -vf scale=480:320,format=i420 -nosound -ovc \
  raw -of rawvideo -ofps 23.976 -o tmp.fifo.yuv \
  max.dv 2>&1 > /dev/null &
x264 -o max-video.mp4 --fps 23.976 --bframes 2 \
  --progress --crf 26 --subme 6 --analyse \
  p8x8.b8x8,i4x4.p4x4 --no-psnr tmp.fifo.yuv 480x320
rm tmp.fifo.yuv
```

Here we instruct mencoder to scale the output to 480x320 and also tell x264 to accept that frame size. This will further reduce the file size, which is appropriate for video on the Web.

Final Steps

Based on the QuickTime format, the .mp4 container format can store many types of media and is also the MPEG standard for storing H.264 video and AAC audio, which is how we will be using it. Use MP4Box,

Need help,
not headaches?

OffAssist

All of your business needs. One easy solution...

BOOKKEEPING
and
VIRTUAL ASSISTANCE

Get your FREE copy of "VA's and the Tech Sector: Plugging a VA into Your Business" at:
www.offassist.com/linuxjournal

www.offassist.com

We Speak Geek! Open Source — Open Mind

Table 1. Table of Results

File	File Size	Video Bitrate
max.dv	32M	3MB/s
max-xvid.avi	623K	418kb/s
max-x264.mp4	522K	392kb/s

which is part of the gpac project, to combine the audio and video streams we've just created:

```
MP4Box -add max-video.mp4 -add audiodump.aac \
-fps 23.976 max-x264.mp4
```

This produces the final output file max-x264.mp4. You can play back the file with MPlayer, or with Apple's QuickTime player on a non-Linux OS. You also can embed this file into a Web page for playback from a browser by using Apple's instructions for embedding QuickTime movies (see Resources). Free software tools such as the mplayer-plugin can be used to play this file from within Firefox on Linux.

By way of comparison, here are the file sizes and bitrates of the original raw DV file max.dv, our H.264-encoded file max-x264.mp4 and a corresponding XviD encoding max-xvid.avi, which was created from the same source video (see Resources):

```
mencoder max.dv -vf scale=480:320 -ovc xvid -xvidencopts \
fixed_quant=7:qpel:nopacked -oac mp3lame \
-ofps 24000/1001 -o max-xvid.avi
```

And below are accompanying screenshots for each sample.



Figure 1. DV



Figure 2. XviD



Figure 3. x264



As you can see, the visual quality of the H.264-encoded file is just as high as the XviD version, arguably higher, but at a lower bitrate and file size. This shows that you can achieve similar results in less space, or much better results in the same space, with H.264 compared to other codecs such as XviD. In addition, the work flow and options for encoding with x264 are similar to XviD, but with greatly improved output. So, if you are used to encoding with XviD, many of the concepts and options should be familiar to you when working with x264.



Figure 4. XviD Detail




Figure 5. x264 Detail

The more you experiment with x264, the more you will discover the amazing savings in bitrates and file sizes while still maintaining an extremely high visual quality. The world of video encoding is definitely a black art, as there are hundreds of variables and options that can be brought to bear in any particular encoding project. There is no one-size-fits-all method of video encoding. However, the technical superiority of H.264 over XviD or regular MPEG-2-encoded video is too great not to take advantage of it. And, you can start taking advantage of it today, using the tools described above. Because H.264 is an MPEG-standard encoding, used with an MPEG-standard audio codec inside of an MPEG-standard container format, all the work you invest in using these tools to encode your video will be future-proof as well as high-quality. Use the techniques outlined above as a starting point for your own H.264-encoding projects, and you'll discover why H.264 is becoming the next standard for video encoding. ■

Resources for this article: www.linuxjournal.com/article/9197.

Dave Berton is a professional programmer. He can be contacted at mosey@freeshell.org.



ASA COMPUTERS

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

Hardware Systems For The Open Source Community—Since 1989

(Linux, FreeBSD, NetBSD, OpenBSD, Solaris, MS, etc.)

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

**AMD Opteron™ Value Server-
\$795**

- 1 U 14.3" Deep
- AMD Opteron 140 1M Cache
- 1 GB DDR ECC REG PC-3200
- 1 of 2 40GB SATA Drive
- 2 X 10/100/1000 NIC
- Options: CD, FD, or Second Drive, Raid
- ADD Your Logo



**iSCSI Dual AMD Opteron™
1U to 8U, Call for Pricing**

- 1TB to 30TB of iSCSI Storage
- Dual AMD Opteron 246
- 1 GB DDR ECC REG PC-3200
- Dual GigE LAN
- Redundant PS, Hot-Swap Drives
- RAID Options, RAID 5, 10, 50
- More Customization is available



**1U SCSI Quad AMD Opteron™ -
Starting @ \$2850**

- 1of 4 AMD Opteron 848
- 2 GB DDR ECC REG PC-3200
- 1 of 3 36GB SCSI Drive
- 2 GigaE, CD, FD,
- Remote Management Card (IPMI)



**30TB AMD Opteron™ Storage
Solution- Starting @ \$26,395**

- 30TB SATA Storage in 8U
- Includes all Raid Cards, Raid 5, 10
- Dual AMD Opteron 246
- 2 GB DDR, ECC REG PC-3200
- Dual GigE, FD, CD



Your Custom Appliance Solution

Let us know your needs, we will get you a solution



"Your Logo Here"

Custom Server, Storage, Cluster, etc. Solutions

Please Contact us for all type of Storage solutions, NAS, DAS, iSCSI, Fiber RAID, SATA, SAS
*Free shipping on selected servers and all notebooks.





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

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

Chapter 10: Personalizing Ubuntu: Getting Everything Just Right

An excerpt from *Beginning Ubuntu Linux: From Novice to Professional* by Keir Thomas, published by Apress. KEIR THOMAS



Beginning Ubuntu Linux: From Novice to Professional
by Keir Thomas
Apress, March 2006
ISBN: 1-59059-627-7
\$39.99 US
608 pages

If you've read this book from Chapter 1, by this stage, you no doubt have become comfortable with Ubuntu. You've started to realize its advantages and are on the way to making it your operating system of choice.

But things might still not be quite right. For instance, you might find the color scheme is not to your tastes. Or perhaps you feel that the mouse cursor moves a little too fast (or too slowly). Maybe you simply want to stamp your own individuality on your system to make it your very own.

That's what this chapter is all about. We look at personalizing Ubuntu so that you're completely happy with your user experience.

Changing the Look and Feel

Ubuntu is similar to Windows in many ways, but the developers behind it introduced improvements and tweaks that many claim make the software easier to use. For example, Ubuntu offers multiple virtual desktops—long considered a very useful user-interface feature that seems to have passed Microsoft by. It also moves the programs menu to the top of the screen, leaving the whole width of the screen at the bottom to display taskbar buttons. This is very sensible, because the buttons don't look cramped when more than a handful of applications are open. However, if you're not satisfied with Ubuntu's out-of-the-box look and feel, you can change it.

You might be used to changing the desktop colors or wallpaper under Windows, but Ubuntu goes to extremes and lets you alter the look and feel of the entire desktop. Everything from the styling of the program windows to the desktop icons can be altered quickly and easily.

Altering the Theme

Ubuntu refers to the look of the desktop as a theme. Because it's built on the GNOME desktop, Ubuntu allows you to radically personalize your desktop theme. Several different themes come with the distribution, and you can download many more themes. Each lets you change the way the windows look, including the buttons and the icon set (although some themes come without additional icons).

However, unlike Windows themes, most GNOME themes don't change the fonts used on the desktop, and the wallpaper and color scheme will probably remain broadly the same. You can change these manually, as described in the "Setting Font Preferences" and "Changing the Wallpaper"

sections a bit later in this chapter.

To alter the theme, select System→Preferences→Theme. Then it's simply a matter of choosing a theme from the list in the Theme Preferences dialog box, as shown in Figure 1. A useful hint is to open a Nautilus file browser window in the background (Places→Desktop), so you can see how the changes will affect a typical window.

Note: The default Ubuntu theme is called Human and is designed to represent the skin tones of the world's population. This is intended to reflect Ubuntu's mission of being accessible to everyone, no matter where or who they are.

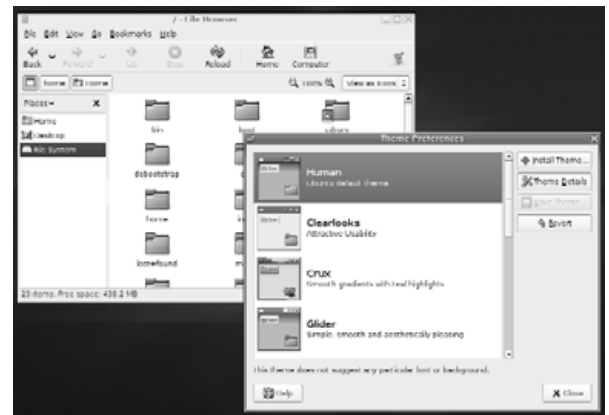


Figure 1. Ubuntu comes with several theme choices.

My favorite themes are Clearlooks and Mist, largely because they're simple and uncomplicated. Remember that you'll be working with the theme on a daily basis, so it should be practical and not too distracting. Those miniature close, minimize, and maximize buttons might look stylish, but they're useless if they're so small that you can't reliably click them with your mouse.

As well as changing the overall theme, you can also modify individual theme components, and even download more theme components.

Changing Individual Theme Components

You can alter the three aspects that constitute a GNOME theme: the controls (sometimes known as widgets), the

window borders, and the icons. Controls are simply the elements you click within dialog boxes and windows: buttons, scroll bars, and so on. The window borders are, as seems obvious, the borders of program windows and dialog boxes, with particular attention paid to the top of the window, where the program name appears along with the minimize, maximize, and close buttons.

Note: To make matters a little confusing, some window borders have their own selection of close, minimize, and maximize controls, which can't be overridden with individual selections for controls.

To make changes to a theme, click the Theme Details button in the Theme Preferences dialog box (Figure 1), and then click each tab to see your choices, as shown in Figure 2. Unfortunately, there are no thumbnail previews of each style, but as soon as you click each option, it will be automatically applied to the currently open windows. To preview the effects fully, the best policy is to keep a Nautilus window open (Places→Desktop).

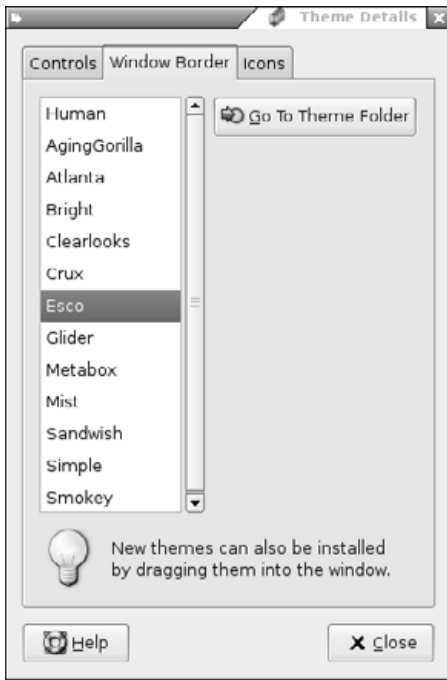


Figure 2. You can create a theme by choosing your own controls, window borders, and icons.

When you've made your choices, you can save the theme for further use. Simply click the Save Theme button in the Theme Preferences dialog box. You'll need to give the theme a name and, if you wish, a short description for future reference. If you don't save the theme, as soon as you select another one, the changes you made will be lost.

www.faircom.com/go/?track

Your packages get

tracked & delivered

on time by FedEx.

FairCom database technology makes it possible.



FairCom[®]

Installing Additional Components

If you get tired of the built-in possibilities, you can download additional theme components, such as window borders and controls, to enhance your desktop experience. You have two ways of getting new themes:

- Download themes from the official Ubuntu repositories.
- Visit the GNOME Art web site (art.gnome.org) and download items from there.

Downloading from Ubuntu Repositories

To get theme components from the Ubuntu software repositories, you use the Synaptic Package Manager. Setting up Synaptic Package Manager to use the online repositories is described in Chapter 8.

Select System→Administration→Synaptic Package Manager, click the Search button, and enter `gtk2-engines` as a search term (`gtk2-engines` is how Ubuntu refers to theme components). In the list of results will be those `gtk2-engines` already installed, indicated by a dark green check box, and several that are available for download.

Icons rarely come in `gtk2-engines` packages, and instead are contained in their own packages. To find icons, use the Synaptic Package Manager to search for `gnome-icon-theme` (without any dashes).

Although each theme component comes with a description, you won't really know what it looks like until you see it. The best policy is to download all of them and audition them one by one. However, be aware that themes can be large, so they may take some time to download on a slower connection.

Caution: Unless you've already installed the KDE desktop, don't download `gtk2-engines-gtk-qt`. This is a piece of system software designed to give KDE applications the same look as GNOME applications. Selecting it will cause the entire KDE desktop subsystem to download, too.

Don't forget that you're downloading theme components, rather than entire themes. To use your new theme components, select System→Preferences→Theme, click the Theme Details button, and choose from the various lists.

Downloading from the GNOME Art Web Site

Visiting the GNOME Art site (art.gnome.org), shown in Figure 3, gives you access to just about every theme ever created for GNOME. In fact, the site also contains wallpaper selections, icons, and much more

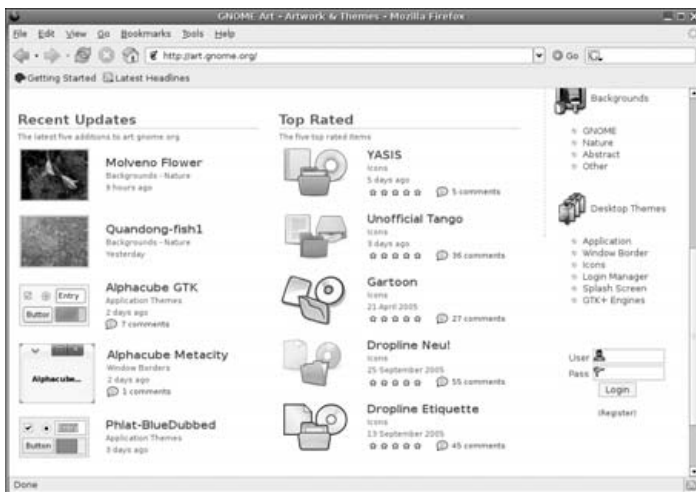


Figure 3. The GNOME Art site contains the latest themes, and you can use all of them with Ubuntu.

besides. All of the offerings are free to use, and most of the packages are created by enthusiasts.

Installing new theme components is easy. If you wish to install a new window border, for example, click the link to browse the examples and then, when you find one you like, click to download it. It will be contained in a `.tar.gz` archive, but you don't need to unpack it. Simply select System→Preferences→Theme and click the Install Theme button in the Theme Preferences dialog box. Then browse to the downloaded theme and click Open. You can delete the downloaded file when you're finished.

Note: The same principle of sharing that underlines the GPL software license is also usually applied to themes. This means that one person can take a theme created by someone else, tweak it, and then release it as a new theme. This ensures constant innovation and improvement.

Changing the Wallpaper

The default Ubuntu wallpaper, Lagoon, is a love-it-or-hate-it affair. Some find its emphasis on dark colors depressing; others appreciate its humanist metaphor. If you're one of those who prefer something different, it's easy to switch. Simply right-click the desktop and click Change Desktop Background. If you want to use a picture of your own as wallpaper, click the Add Wallpaper button, and then browse to its location.

In the Style drop-down list, you can select from the following choices:

- Centered: This option places the wallpaper in the center of the screen. If the wallpaper is not big enough to fill the screen, a border appears around the edge. If it's bigger than the screen, the edges of the wallpaper are cropped off.
- Fill Screen: This option forces the picture to fit the screen, including squashing or expanding it if necessary (known as altering its aspect ratio). If the wallpaper isn't in the same ratio as the screen, it will look distorted. Most digital camera shots should be okay, because they use the same 4:3 ratio as most monitors (although if you have a wide-screen monitor, a digital camera picture will be stretched horizontally).
- Scaled: Like the Fill Screen option, this option enlarges the image if it's too small or shrinks it if it's too big, but it maintains the aspect ratio, thus avoiding distortion. However, if the picture is in a different aspect ratio than the monitor, it may have borders at the edges.
- Tiled: If the picture is smaller than the desktop resolution, this option simply repeats the picture (starting from the top left) until the screen is filled. This option is primarily designed for patterned graphics.

Don't forget that the GNOME Art web site (art.gnome.org) offers many wallpaper packages for download.

Tip: Looking for some good wallpaper? Visit Flickr (www.flickr.com). This is a community photography site where many people make their pictures publicly available.

Setting Font Preferences

Ubuntu lets you change the fonts that are used throughout Ubuntu (referred to as system fonts). You can also alter how they're displayed.

To change a system font, select System→Preferences→Font. In the Font Preferences dialog box, shown in Figure 4, click the button next to the system font you want to change, and then choose from the list. You can also set the font point size so, for example, you can make the

Advertiser Index

For advertising information, please contact our sales department at 206-782-7733 ext. 2 or ads@ssc.com.
www.linuxjournal.com/advertising

labels beneath icons easier to read.

By clicking the entries under the Font Rendering heading in the Font Preferences dialog box, you can change how fonts look on your monitor. This will alter the antialiasing and hinting of the font. Antialiasing softens the edges of each letter to make them appear less jagged. Hinting affects the spacing and shaping of the letters. Used together, they can make the on-screen text look more pleasant. Try each Font Rendering setting in sequence to see which looks best to you (the text in the dialog box will update automatically to show the changes). Nearly everyone with a TFT-based screen, including notebook users, finds the Subpixel Smoothing option best.



Figure 4. You can alter the way fonts appear on screen by using the Fonts Preferences dialog box.

Bytecode Hinting

Two font hinting subsystems are available under Ubuntu: Autohinting and Bytecode Interpreting. There's a lengthy debate about which produces the best results. Personally, I prefer to use the Bytecode Interpreter because I believe it leads to the cleaner fonts, but others say Autohinting is better in this regard.

Unfortunately, the Bytecode Interpreter is protected by patents in the United States (for more information, see www.freetype.org/patents.html). This means that individuals in the United States can't use it without paying a license fee to the patent holder (in theory, at least). Those in the rest of world should be fine.

It's easy to activate the Bytecode Interpreter. Follow these steps:

1. Open a GNOME Terminal window (Applications→Accessories→Terminal).
2. In the terminal window, type the following:

```
sudo dpkg-reconfigure fontconfig
```

3. On the first screen of the configuration program, select Native to activate the Bytecode Interpreter. Alternatively, you can choose Autohinting or even None, which will turn off the hinting system.
4. On the next screen, you can select whether subpixel rendering is activated. This is useful only for TFT screens (including notebooks), so either select Automatic or, if you use a TFT monitor and want to

Advertiser	Page #	Advertiser	Page #
ABERDEEN, LLC www.aberdeeninc.com	47	LINUXWORLD UK www.linuxworldexpo.co.uk	95
AML www.amltd.com	65	LOGIC SUPPLY INC www.logicsupply.com	85
APACHECON www.apachecon.com	57	LPI www.lpi.org	63
APPRO HPC SOLUTIONS appro.com	C2	MBX www.mbx.com	21
ASA COMPUTERS www.asacomputers.com	35, 77	MICROWAY, INC. www.microway.com	C4
AVOCENT CORPORATION www.avocent.com/special	1	MIKRO TIK www.routerboard.com	7
CARL.NET www.carl.net	93	MONARCH COMPUTERS www.monarchcomputer.com	10, 11
CIARA TECHNOLOGY www.ciaara-tech.com	8, 9	OFFASSIST www.offassist.com	75
CORAD, INC. www.coraid.com	19	OPEN SOURCE STORAGE www.opensourcestorage.com	13
COYOTE POINT www.coyotepoint.com	3	OPERA SOFTWARE www.opera.com	5
EMAC, INC. www.emacinc.com	88	PENGUIN COMPUTING www.penguincomputing.com	25
EMPERORLINUX www.emperorlinux.com	15	POLYWELL COMPUTERS, INC. www.polywell.com	55
FAIRCOM CORPORATION www.faircom.com	79	THE PORTLAND GROUP www.pgroup.com	27
FREE SOFTWARE FOUNDATION member.fsf.org	94	RACKSPACE MANAGED HOSTING www.rackspace.com	C3
GECAD TECHNOLOGIES/AXIGEN www.axigen.com	23	R CUBED TECHNOLOGIES www.rcubedtech.com	71
GENSTOR SYSTEMS, INC. www.genstor.com	83	SERVERS DIRECT www.serversdirect.com	43
HURRICANE ELECTRIC www.he.net	39	SILICON MECHANICS www.siliconmechanics.com	41, 61
INTEL www.intel.com/software/products	45	SUPERMICRO www.supermicro.com	37
IRON SYSTEMS www.ironsystems.com	91	TECHNOLOGIC SYSTEMS www.embeddedx86.com	46
JTL NETWORKS www.jtl.net	6	THINKMATE www.thinkmate.com	17
LANTRONIX www.lantronix.com	31	USENIX LISA www.usenix.org/events/lisa06	89
LAYER 42 NETWORKS www.layer42.net	87	ZT GROUP INTERNATIONAL www.ztgroup.com	33

ensure the option is activated, select Always.

5. The third screen offers the option of using bitmap fonts. These are fonts that, unlike the TrueType fonts used within the rest of Ubuntu, don't scale beyond their original size. There's no harm in enabling them because they can sometimes be useful as system fonts.
6. Once the program has finished configuring the software, restart your X server by logging out and then back in again.

Configuring Input Devices

Mouse and key repeat speeds are personal to each user, and you may find the default Ubuntu settings not to your taste, particularly if you have a high-resolution mouse such as a gaming model. Fortunately, changing each setting is easy. You'll find the relevant options under the System→Preferences menu.

Configuring Mouse Options

Select System→Preferences→Mouse to open the Mouse Preferences dialog box, which has three tabs:

- **Buttons:** This tab lets you set whether the mouse is to be used by a left-handed or right-handed person. Effectively, it swaps the functions of the right and left buttons. Beneath this is the double-click timeout setting. This is ideal for people who are less physically dexterous because the double-click speed can be slowed down. On the other hand, if you find yourself accidentally double-clicking items, you can speed it up.
- **Cursors:** On this tab, you can select from any mouse cursor themes that are installed. You can also activate the Locate Pointer option, which causes a box to appear around the mouse cursor when you press the Ctrl key. This can help you find the cursor on a busy desktop.
- **Motion:** This tab, shown in Figure 5, lets you alter the speed of the mouse pointer, as well as the drag-and-drop threshold. Changes are made as each setting is adjusted, so to test the new settings, simply move your mouse. Here's what the settings do:
 - › The Acceleration setting controls how fast the mouse moves. Whenever you move the mouse, the pointer on screen moves a corresponding amount. However, the cursor actually increases in speed the more you move your hand (otherwise, you would need to drag your hand across the desk to get from one side of the screen to the other). This is referred to as acceleration. If you set the acceleration too high, the pointer will fly around the screen, seemingly unable to stop. If you set it too slow, you'll need to ramp the mouse several times to make it go anywhere.
 - › The Sensitivity setting controls how quickly the acceleration kicks in when you first move the mouse. Choosing a higher setting means that you can move the mouse relatively quickly before it starts to accelerate and cover more screen space. A low setting means that acceleration will begin almost as soon as you move the mouse. Higher sensitivity settings give you more control over the mouse, which can be useful if you use image-editing programs, for example.
 - › The Threshold setting determines the amount of mouse movement allowed in a click-and-drag maneuver before the item under the cursor is moved. This setting is designed for people who have limited dexterity and who might be unable to keep the

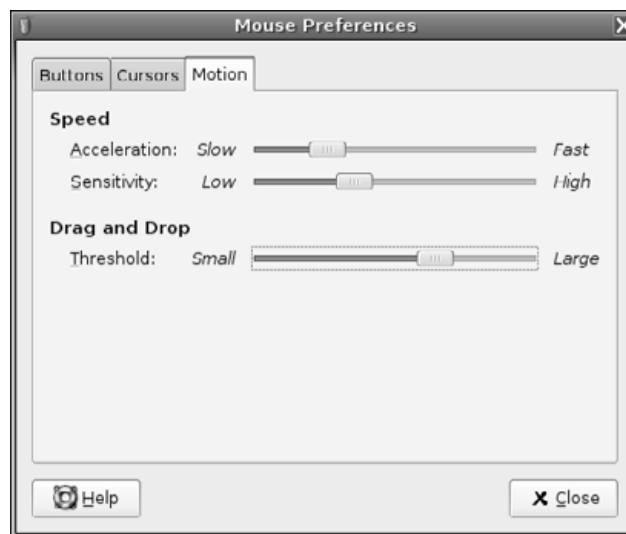


Figure 5. The Mouse Preferences dialog box lets you tame that mouse.

mouse perfectly still when clicking or double-clicking an icon. In such cases, a large threshold value may be preferred.

Changing Keyboard Settings

Select System→Preferences→Keyboard to open the Keyboard Preferences dialog box. This dialog box has four tabs:

- **Keyboard:** Using the Keyboard tab, you can alter the rate of key repeat. This can be useful if you often find yourself holding down the Backspace key to delete a sentence; a shorter setting on the Delay slider and a faster setting on the Speed slider can help. However, if you get the settings wrong, you may find double characters creeping into your documents; typing an f may result in ff, for example. Beneath the Repeat Keys setting is the Cursor Blinking slider. Altering this may help if you sometimes lose the cursor in a document; a faster speed will mean that the cursor spends less time being invisible between flashes.
- **Layouts:** On the Layouts tab, you can choose to add an alternative keyboard layout, as shown in Figure 6. For example, if you write in two different languages on your keyboard, it may be helpful to be able to switch between them. Click the Add button and select the second language from the list.
- **Layout Options:** This tab lets you select from a variety of handy tweaks that affect how the keyboard works. For example, you can configure the Caps Lock key to act like a simple Shift key, or you can turn it off altogether. You can configure the Windows key so that it performs a different function, too. Put a check alongside the option you want after reading through the extensive list of options.
- **Typing Break:** This tab features a function that can force you to stop typing after a predetermined number of minutes. It does this by blanking the screen and displaying a "Take a break!" message. Note that a notification area icon will appear before the break time to give you advanced warning of the lockout.

Creating Keyboard Shortcuts

Ubuntu lets you define your own keyboard shortcuts for just about any action on the system. To create a shortcut, select System→Preferences→Keyboard Shortcuts. In the dialog box,

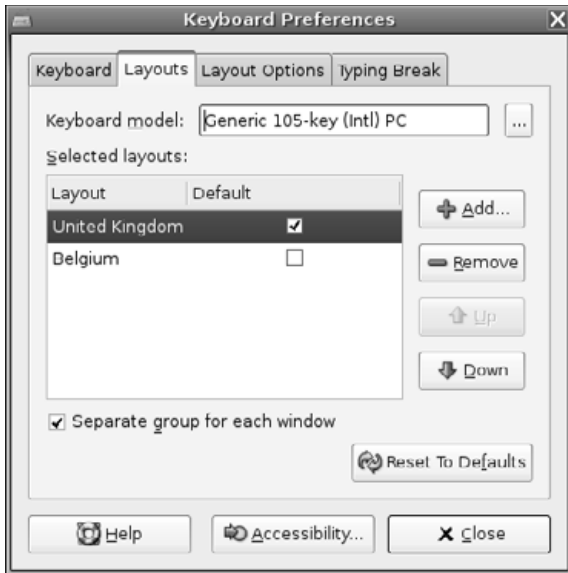


Figure 6. You can have more than one language setting in place for a keyboard, which is handy if you need to type in a foreign language.

search through the list for the action you want to create a shortcut for, click it, and then press the key (or key combination) you want to use. For example, you might locate the Volume Up and Volume Down entries in the list, click each, and then press Ctrl+left arrow and Ctrl+right arrow. Then you will be able to turn the volume of your sound card up or down by holding down Ctrl and tapping the left or right arrow key, respectively.

Caution: Be careful not to assign a shortcut to a popular key. It might be nice to make Totem Media Player appear when you hit the spacebar, for example, but that will mean that it will start up several times whenever you type a sentence in a word processor! Also be aware that some key combinations are used by applications. Within OpenOffice.org's Writer, for example, the Ctrl+left/right arrow key combination moves you from word to word in a paragraph. If you define those combinations as shortcuts, you will no longer have this functionality.

I like to configure my /home folder to appear whenever I press the Home button on the keyboard. This can be done by locating the Home Folder option under the Desktop heading.

Personalizing Login Options

You can even personalize the login screen under Ubuntu. This is known technically as the GNOME Display Manager, or GDM. To access its configuration options, select System→Administration→Login Screen Setup. The dialog box has eight tabs:

- **General:** In the Local drop-down list, you can choose the type of login screen: Themed Greeter, which is to say one that includes the Ubuntu graphic, or GTK+ Greeter, which is a basic login box into which you can type your details. By unchecking the Use Default Welcome boxes, you can type your own Welcome and Remote Welcome text (the remote text is what appears if someone logs in graphically via GDM across a network; it isn't something you should worry about). The Automatic Login heading lets you do away with the login box completely and go straight to the desktop from bootup. Simply put a check in the box and provide the login username. This presents obvious security issues, but if you're the only person using



Company With a Vision

Great Minds Great Solutions

1-877-25-SERVER

www.genstor.com

Customized Solutions for

SERVERS :: **STORAGE** :: **APPLIANCES**

Linux - FreeBSD - x86 Solaris - MS etc.

SERVERS

LOW POWER – BUDGET – HIGH DENSITY

(1U, 2U, 3U and above - Xeon, P4, Opteron, VIA)



2U Dual Xeon/Opteron Server
Upto 24GB RAM
Upt o 8 SATA/SCSI HDD
Starting @ \$ 2000.00
Call for your custom needs

SERVER MANAGEMENT & MORE

Genstor offers Data Center-proven Linux Management Technology - which gets you :



Provisioning
Deployment
Change Management for Linux
Environments large and small
Track Changes
Instant Rollback and much more

Setup your racks with heterogeneous Linux environments in minutes

Contact sales@genstor.com

STORAGE

SATA- NAS-DAS-iSCSI -SAS Storage Solutions



5U Dual Xeon/
Opteron SATA Storage
Upto 24GB RAM
Upto 26 SATA HDD
Upto 18TB Storage
Starting @ \$ 3930.00
Call for your custom needs

CUSTOM APPLIANCE SOLUTIONS

Prototype – Certifications - Logo Screen Printing



Custom Turnkey OEM Appliance Solutions
From Prototype to Drop Ship
Custom OS/Software Image Installs
No Quantity is small to Customize
Call for your Custom OEM Solutions

FACING POWER PROBLEMS AT DATA CENTER PL. CALL FOR SOLUTIONS

Contact Genstor For Your Hardware Needs

Genstor specializes in customizing hardware around the OS of your choice (Linux, *BSD, x86 Solaris etc.) with Intel and AMD offerings. Please contact Genstor sales@genstor.com for all your hardware needs.

GENSTOR SYSTEMS, INC.

780 Montague Exp. # 604, San Jose, CA 95131

Phone: 1-877-25 SERVER or 1-408-383-0120

Email: sales@genstor.com

Fax: 1-408-383-0121

Prices, Products and Availability subject to change without notice.

the computer and if it's located in a secure location, you might want to choose this option. The Timed Login option lets you select a user who will be logged in by default after a given period. This is useful if you want to present the opportunity to log in as a different user but also want to have the failsafe of logging in automatically, too.

- **GTK+ Greeter:** Here, you can change the appearance of the simple GTK+ Greeter that will appear if you chose the relevant option on the General tab. You can change the picture by clicking the Browse button, or get rid of it completely by clicking No Logo. You can also alter the background so that there's a picture or color.
- **Themed Greeter:** The Themed Greeter is the default under Ubuntu, and it makes the standard login page appear. You can select from a couple of other themes from the list on the left. Particularly handy is the Happy GNOME with Browser option, which shows a list of all users on your system during login. To log in, just click the appropriate entry and type the necessary password. Clicking the Install New Theme button lets you install a Themed Greeter theme from disk. You can download these from art.gnome.org.
- **Security:** This tab lets you alter login settings that might present a security risk to your system. For example, if the Happy GNOME with Browser option is activated in the Themed Greeter tab, you can activate the Face Browser, which will show a picture on the login screen. This is considered a security risk, because it removes the need for people to type in their usernames, thus potentially handing valuable information to hackers. You can also activate the X Display Manager Control Protocol (XDMCP) for GDM, which will let users log in graphically using GDM from a remote computer (this requires the remote system to be specially configured).
- **XServer:** This tab lets you specify settings that can be passed to the X server when it's started by GDM. In most cases, you can ignore this tab.
- **Accessibility:** Here, you can allow activation of the GNOME Assistive Technology modules. You can also have GDM make sounds when various events occur, such as a bad login. This can help partially sighted people.
- **XDMCP:** If you activated XDMCP for GDM in the Security tab, this tab lets you fine-tune various settings, such as the network port it listens on as well as the maximum number of remote sessions. Unless you have specifically activated XDMCP, you can ignore this tab.
- **Users:** Here, you can specify which users are offered as choices within GDM if the Face Browser option is activated in the Security tab. Bear in mind that Linux has many system user accounts that aren't designed to allow logins. By default, all users who have a password are displayed, which is the best way of working (the system accounts don't have passwords because they aren't login accounts).

Setting Up Assistive Technology Support

You might know about the Accessibility tools under Windows, which help people with special needs use the computer. It's possible to use an on-screen magnifier so that users can better see what they're typing or reading, for example.

Under the GNOME desktop, the Accessibility tools are referred to as Assistive Technology Support. To use them, you need to install additional software packages, and then enable them. Follow these steps:

1. Assuming the Synaptic Package Manager is set up to use the online repositories (see Chapter 8), open the program (System→Administration→Synaptic Package Manager).
2. Click the Search button and enter gok as a search term. In the list of results, click the check box alongside the gok entry, and then click Mark for Installation.
3. Click the Search button again and search for gnompernicus. Again, mark its entry for installation. Then click Apply.
4. Once the packages are installed, select System→Administration→Preferences→Assistive Technology Support. Click the check box alongside Enable Assistive Technologies.
5. Choose from the list the features you would like to use. They will then start automatically the next time you log in. The options work as follows:
 - › The Screenreader uses a speech synthesizer to announce whatever you click on, as well as whatever you type. To alter its settings, click the Preferences button in the Gnompernicus dialog box, and then click the Speech button in the Preferences dialog box.
 - › The GNOME Onscreen Keyboard (GOK) can be used by a mouse, but is most useful when an alternative input device is used, such as a touch screen. As well as presenting a virtual keyboard, it shows the options on screen as a large and easy-to-activate series of buttons. For more information, click the Help button when GOK starts.
 - › The Magnifier divides the screen into two halves. The right side displays a magnified version of the left side. To learn more, click the Help button in the Gnompernicus dialog box.
 - › Gnompernicus also includes support for Braille output devices. To learn more, click the Help button.

Changing Your Login Picture

If you activate the Face Browser feature along with the Happy GNOME with Browser option, GDM will display a picture alongside your username on the login screen, as shown in Figure 7. You can then click this and type your password to log in. You might be familiar with a similar system under Windows XP.

You can choose your login picture by clicking System→Preferences→Login Photo. Ubuntu offers several pictures to choose from, or you can click the Browse button to locate your own. Ideally, the image you choose should be square and 96x96 pixels, although if the picture is too large, it will be automatically scaled down.

Adding and Removing Desktop Items

Virtually the entire Ubuntu desktop can be redesigned and restructured. You can move the Applications menu from the top of the screen to the bottom to be more like Windows, for example, or you can add numerous desktop shortcuts to popular applications and/or files.

Adding a Shortcut

Ubuntu's nearest equivalent to a Windows-style desktop shortcut is a launcher. An important difference, however, is that launchers are designed to run a certain command. Therefore, they can only point at programs (although you could create a launcher that contained a command chain required to run a particular program and file; to



Figure 7. The Face Browser lets each user choose an icon to appear on the login screen.

use The GIMP to open a picture, you might create a launcher that ran `gimp picture.jpg`, for example).

If you want to make a shortcut to a data file, such as a picture, you need to create a link. This is just as easy as creating a launcher.

Creating a Launcher

You can create a launcher two ways. One way is to simply click and drag an icon from one of the main menus to the desktop. This effectively copies the menu's launcher to the desktop, rather than creating a new launcher, but the effect is the same.

Note: Not all menu items can be dragged and dropped to create



Figure 8. Creating a launcher is easy. Just fill in the Name and Command fields, and choose an icon.

launchers. In particular, most items on the System menu cannot be dragged and dropped.

The other way to create a launcher is to right-click the desktop and select Create Launcher. In the Create Launcher dialog box, you need to fill in only the Name and Command field; the others fields can be left blank. The Command field must contain a Linux executable program, command, or script. If you use a command-line program or script, you must check the Run in Terminal box. This will open a GNOME Terminal window automatically and run the command or script within it. The terminal window will disappear as soon as the command has finished.

To choose an icon for your launcher, click the Icon button in the Create Launcher dialog box. If you don't choose an icon, the stock GNOME one is used (the same icon as is used for unidentified and/or system files). You can select from several predefined icons or choose your own picture by clicking the Browse button, as shown in Figure 8.

Creating a Link

Although you can create a link to a program, which will then work the same as a launcher, links are usually used to create shortcuts to files. If you then double-click the shortcut, the application associated with that file type will open the file. If you create a shortcut to a picture, for example, when you double-click it, the Eye of GNOME previewer will start.

To create a link, locate the file you want to create the link to, right-click it, and select Make Link. Then copy the new link to wherever you want it to appear, such as the desktop. You don't need to choose an

Logic Supply offers a full line of Mini-ITX mainboards and systems for embedded applications.



Mini- and Nano-ITX mainboards

with VIA, Pentium, or Pentium M processors

- x86 compatibility
- full-featured SBC
- off-the-shelf availability



fanless Mini-ITX systems

Utilizing heat pipe technology, these completely fanless Mini-ITX systems offer the perfect solution for applied computing and embedded environments.

LEADERS IN MINI-ITX SOLUTIONS

LOGIC
SUPPLY

LOGICSUPPLY.COM
info@logicsupply.com
CALL US AT 802 244 8302

icon, because the link inherits the icon of the original file. For example, if it's a picture link, it will inherit the thumbnail preview icon.

Note: If you find the Make Link option grayed out, it's likely that you don't have sufficient permissions to write the link to the directory in question.

Adding and Removing Menus

You can add either the Applications menu or the entire set of menus (Applications, Places, and System) to the bottom of the screen. This can help those who long for the Windows Start button approach to access programs.

Adding All the Menus to a Panel

To add the Application, Places, and System menu to the Panel at the bottom of the Ubuntu desktop, follow these steps:

1. Right-click a blank spot on the bottom Panel and select Add to Panel.
2. In the dialog box that appears, click the Menu Bar option to add all three menus.
3. Click the Add button at the bottom of the dialog box.

Adding the Applications Menu to a Panel

When you add just the Applications menu, it appears as a single icon. To move it to the bottom-left side of the desktop, in the same location as the Windows Start menu, you need to delete and then re-create the Show Desktop icon and the Window List. Here are the steps:

1. Right-click a blank spot on the bottom Panel and select Add to Panel.
2. In the dialog box that appears, click the Main Menu option.



Figure 9. If you just can't do without that Start button, you can re-create one on your Ubuntu desktop.

3. Click the Add button at the bottom of the dialog box.
4. Right-click the Window List's grab bar and select Delete.
5. Right-click the Show Desktop icon and select Delete.
6. Right-click the new menu icon, select Move, and relocate it to the bottom left.
7. Right-click the Panel and select Add to Panel.
8. In the Desktop & Windows category, select Window List, and then click Add.
9. Repeat steps 7–8 to add the Show Desktop button.

When you're finished, the Applications menu should appear as shown in Figure 9.

Deleting a Menu

Creating new instances of the menus won't delete the old ones. If you create a new Applications menu at the bottom of the screen, for example, the old Applications menu will remain at the top of the screen. In fact, you can have as many instances of the menus on the desktop as you wish, although this won't be a good use of desktop space!

To delete any menu, simply right-click anywhere on that menu and select Remove from Panel.

Give Me My Trash Can!

The developers who designed Ubuntu's desktop decided to keep the desktop clean of icons. This included relegating the Wastebasket icon to its own applet at the bottom-right side of the screen. Many people find using the applet a little difficult and miss the desktop trash can icon, which has been present on Windows and Mac OS desktops for more than 20 years.

The good news is that it's easy to get the trash can back. Click Applications→System Tools→Configuration Editor. In the program window that appears, click the down arrows next to Apps, then Nautilus, and then Desktop. On the right side of the program window, put a check in the trash_icon_visible entry.

Alternatively, in the Configuration Editor, click Edit→Find and enter trash_icon_visible as a search term. Make sure that the Search Also In Key Names box has a check in it. Then click Find. The results will be listed at the bottom of the program window. Click the /apps/nautilus/desktop/trash_icon_visible entry. Then make sure there's a check in the trash_icon_visible box.

Be careful when using the Configuration Editor program. It lets you configure just about every aspect of the GNOME desktop and doesn't warn you when you're about to do something devastating, so the potential for accidental damage is high!

Personalizing the Panels

Panels are the long strips that appear at the top and bottom of the Ubuntu screen and play host to a choice of applets and/or icons. You can add a new panel by simply right-clicking an existing one and selecting New Panel, or you can remove a panel by right-clicking it and selecting Delete This Panel.

Caution: If you delete a panel, the arrangement of applets it contains will be lost. Of course, you can always

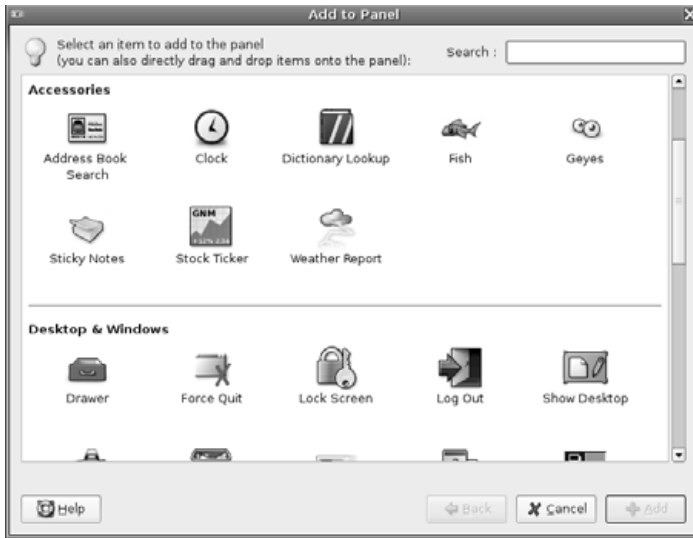


Figure 10. A wide variety of applets are available. Some are informative; others are just fun.

re-create the collection of applets on a different panel.

By right-clicking a panel and selecting Properties, you can change its size and dimensions. For example, by unchecking the Expand box, you can make the panel shrink to its smallest possible size. Then, when you add new applets (or, in the case of a panel containing the Window List, a new program is run), the panel will expand as necessary. This can be a neat effect and also creates more desktop space.

Selecting the Autohide feature will make the panel slide off the screen when there isn't a mouse over it. Choosing Show Hide Buttons will make small arrows appear on either side of the panel so that you can click to slide it off the side of the screen when it's not in use.

Working with Applets

Almost everything you see on the desktop is an applet, with the exception of shortcut icons and the panels. A menu is a form of applet, for example, as is the Workspace Switcher.

Ubuntu provides many more applets that you can choose to add to the desktop to provide a host of useful or entertaining functionality. To add an applet, right-click the Panel and select Add to Panel. As shown in Figure 10, you have a wide choice of applets, divided into categories. Many require configuration when they've been created, so you may need to right-click them and select Properties. For example, you'll need to set your location in the Weather Report applet's properties so it can provide accurate forecasting.

To remove an applet, simply right-click it and select Remove from Panel.

Setting Power-Saving Features

Ubuntu includes a number of features to use the power-saving features of your computer, including the ability to switch off the monitor after a set period of inactivity. However, some quick configuration is necessary to set up the system just the way you want it.

Tip: If your computer has a CPU that can adjust its clock speed on the fly, such as a mobile processor or an AMD chip with the PowerNow! function, Ubuntu will automatically install software that will make this work. To see the speed of your processor, right-click the Panel, select Add to Panel, and choose the CPU Frequency Scaling Monitor.

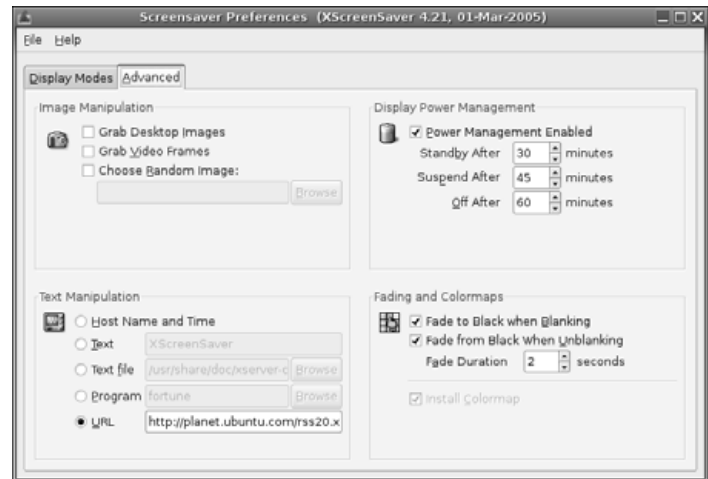


Figure 11. You can configure your monitor to automatically switch itself off after a certain period.

Controlling the Monitor Shutoff

You can configure the monitor to go into standby mode after a certain amount of time has passed. This can save a lot of

LAYER 42™

- Redundant UPS and generator
- Nationwide network
- Free tech support

2U 256kbps ~80GB \$60/mo.	4U or Mid-tower 256kbps ~80GB \$80/mo.
1/4 Rack 512kbps (14U) ~165GB \$200/mo.	1/2 Rack 1mbps (28U) ~330GB \$350/mo.

www.layer42.net

All prices include 100Mbps port, Firewall,
24x7 Monitoring and DNS hosting

408-450-5740
2312 Walsh Ave., Santa Clara, CA 95051

electricity should you happen to leave your computer unattended for long periods.

To configure this aspect of Ubuntu, select System→Preferences→Screensaver, and click the Advanced tab, as shown in Figure 11. In the Display Power Management section, make sure the Power Management box is checked. Set the desired timings in the Standby, Suspend, and Off boxes. Not all monitors support all three features. Consult your monitor's manual to determine which you should configure. If you're unable to find this information, simply set all three to similar values. For example, set Standby to 30 minutes, Suspend to 31, and Off to 32.

Spinning Down the Hard Disk

All modern hard disks come with the ability to spin down their motors to save energy. Then, when data is requested, the motors spin up again. There may be a slight delay while this happens, and some people dislike using disk spin-down because of this. However, on a notebook, it can lead to a substantial increase in battery life. On a desktop system, it's worth considering, because over the lifetime of a computer, it can save a lot of electricity (and therefore money!).

The spin-down settings are contained in the `/etc/hdparm.conf` file, which you'll need to edit by hand. Follow these steps to adjust the spin-down settings:

1. Open a GNOME Terminal window (Applications→Accessories→Terminal).

2. Type the following in the terminal window:

```
sudo gedit /etc/hdparm.conf
```

3. Click Search→Find and, in the box, type `spindown_time`.
4. Click the Find button, and then close the Search dialog box.
5. Change the line the Find routine has highlighted to remove the hash mark from the beginning, so it reads like this: `spindown_time = 24`. You can alter the value to anything you want. Each time unit is five seconds, so 24 equates to 120 seconds (24x5 seconds), or 2 minutes. To set a time of 20 minutes, enter 240 (240x5 seconds). If you specify a number above 240, the time units are increased to 30 minutes. In other words, a value of 241 will equate to 30 minutes, a value of 242 will equate to 60 minutes, and so on.
6. When you've finished, save the file.
7. Reboot for the settings to take effect.

Power Saving—Is It Worth It?

An average computer draws anywhere between 100 to 500 watts of power. An average light bulb draws around 150 watts of power, so you can see that, relatively speaking, computers are low power consumers compared to many household devices. However, it's still worth considering employing power-saving techniques. You might not save yourself a lot of money, but if you switch on power saving, and your neighbor does too, and her neighbor does, then the cumulative effect will add up, and we can all contribute less towards global warming.

Try to avoid leaving your computer turned on overnight, or when you're away from it for long periods. As well as saving power, switching off your computer will avoid wear and tear on its components. Although the CPU can work 24/7 without trouble, it's cooled by a fan that's a simple mechanical device. There are other fans in your computer too, such as the graphics card fan and case fan. Each of these will eventually wear out. If your graphics card fan stops working, the card itself will overheat and might burn out. The same is true of the CPU fan. However, by shutting down your computer overnight, you can effectively double the life of the fans and radically reduce the risk of catastrophic failure. Isn't that worth considering?

Summary

In this chapter, you've learned how to completely personalize Ubuntu to your own tastes. We've looked at changing the theme so that the desktop has a new appearance, and we've examined how to make the input devices behave exactly as you would like.

In addition, you've learned how to add and remove applets from the desktop in order to add functionality or simply make Ubuntu work the way you would like.

Finally, we looked at the power-saving functions under Ubuntu and how you can avoid your computer wasting energy.

In the next chapter, we will look at what programs are available under Ubuntu to replace those Windows favorites you might miss. ■

EMBEDDED LINUX

Made Easy!

PC Compatible
SBCs



Panel PCs



SoM & PC/104
Modules



Custom
Engineering





Embedded
Servers



Turn-Key Solutions!

- ECLIPSE IDE with full Library and Module Support
- Integrated GDB Source Level Debugger
- Resident Menu Driven Setup Utility
- No Charge for Tools or Standard Linux Install
- Free Integration, Setup and Test

Since 1985
OVER
21
YEARS OF
SINGLE BOARD
SOLUTIONS



2.4 Kernel

EMAC, inc.

Phone 618 - 529 - 4525 Fax 618 - 457 - 0110
2390 EMAC Way, Carbondale, Illinois 62902
World Wide Web: <http://www.emacinc.com>

Sponsored by

THE USENIX SIG FOR
USENIX & [sage]
SYSADMINS

6 DAYS OF TRAINING
BY INDUSTRY EXPERTS, INCLUDING:

- Gerald Carter on Ethereal and the Art of Debugging Networks
- Richard Bejtlich on TCP/IP Weapons
- Eelen Frisch on Administering Linux in Production Environments
- Chip Salzenberg on Higher-Order Perl
- And 55 other tutorials

LISA '06

20TH LARGE INSTALLATION
SYSTEM ADMINISTRATION CONFERENCE

A **Blueprint** for Real World
System Administration

DECEMBER 3-8, 2006 | WASHINGTON, D.C.

3-DAY TECHNICAL PROGRAM

Keynote: Cory Doctorow, science fiction writer, co-editor of Boing Boing, and former Director of European Affairs for the EFF, on Hollywood's Secret War on Your NOC

20+ Invited Talks, including:

- Simple Nomad, Vernier Networks, Inc., "Corporate Security: A Hacker Perspective"
- DJ Byrne, Jet Propulsion Laboratory, "Open Source Software and Its Role in Space Exploration"
- Mazda Marvasti, Integrien, "Everything You Know About Monitoring Is Wrong"

Refereed Papers, Hit the Ground Running Track, Guru Is In Sessions, Vendor Exhibition, Workshops, BoFs, WiPs, and more!

Register by November 10 and save! www.usenix.org/lisa06/lja

LISA '06 offers the most in-depth, real-world system administration training available!

Digital Photography and Linux

A primer for various digital photography programs available for Linux.

ADRIAN KLAVER

This article presents a basic overview of some of the programs available for Linux to create and work with digital images. The programs discussed here are those that were installed by default or by the package management programs of SUSE and Kubuntu on my systems. They are fairly standard and should be available for the major distributions. My preferred desktop environment is KDE, so the choices are slanted in that direction, though I indicate where a program also works on GNOME, or where there is an alternative. This is becoming less of an issue, as it is possible to run GNOME applications on KDE and vice versa. In fact, as the screenshots show, I did all my work in KDE. To bring some organization to the article, the programs are introduced in work-flow order. See the on-line Resources for versions of software used.

Scanning with Kooka

To convert images in a nondigital format, it is necessary to scan the image and capture the output into a digital file. The software of choice for this is SANE (Scanner Access Now Easy). SANE takes care of connecting to the physical devices (the scanners) and passing information to and from them. Any one of a number of front ends handles the user interface for the scanner. The two covered here are Kooka and XSane. Kooka is an official part of the KDE Graphics Package, and as such, is found in the KDE menu. XSane is an independent project. Support for scanners is provided by the libraries from the SANE Project, so if you need to determine whether your scanner will work, go to the SANE Web site (see Resources). To scan, open Kooka from the KDE Graphics menu item (Figure 1).

Scanning a color photograph entails selecting a Scan mode of color

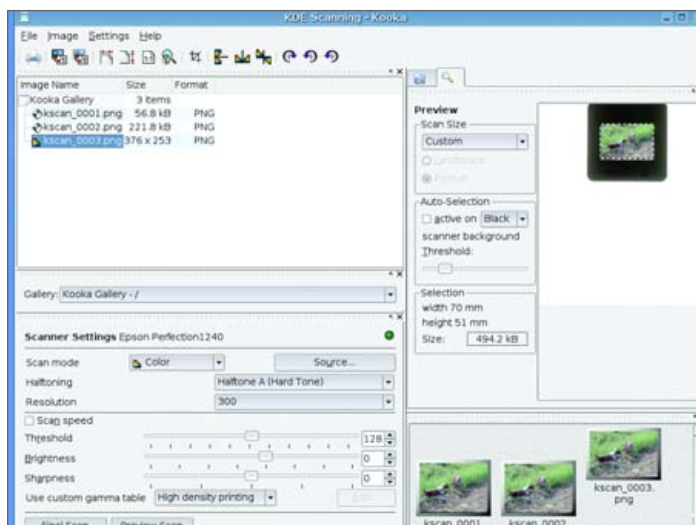


Figure 1. Kooka Showing a Preview with an Area Selected for the Final Scan

and the appropriate resolution. The Source item also may need to be selected (for instance, I use a Epson 1240U that has the choice of a flatbed scanner or transparency unit). Select Preview Scan. This generates an image in the preview box to the right. It is then possible to use your mouse to select the portion of the preview that you want to scan. Select Final Scan. At the end of the scanning process, a dialog box pops up, asking you in which format to save the scanned image. Saving the image creates a thumbnail image in the box at the bottom right and adds the image to the Kooka Gallery at the top left. For further work with an image, select the Image menu item and then Open Graphic Application to use the image editor you prefer.

Scanning with XSane

XSane can be run either in KDE or GNOME. To open XSane, again go to the Graphics menu item and select XSane. One immediate difference you will notice is that XSane opens a host of windows (Figure 2). I generally have the Standard Options, Advanced Options and Preview windows open in addition to the Program window. This can be controlled by clicking on the Window menu item in the Program window and selecting or deselecting the appropriate items.

The information in the Advanced Options window will vary according to the capabilities of the scanner you are using. To start a scan, go to the Program window and select the necessary values. To scan a print, select Viewer, Color, flatbed, full color range and 300 in descending order for the output, image type, source, color range and resolution choices. If the output is going to an inkjet, in the Advanced Options section, set the Color correction to Inkjet Printers.

With the settings taken care of, click the Acquire preview button in the

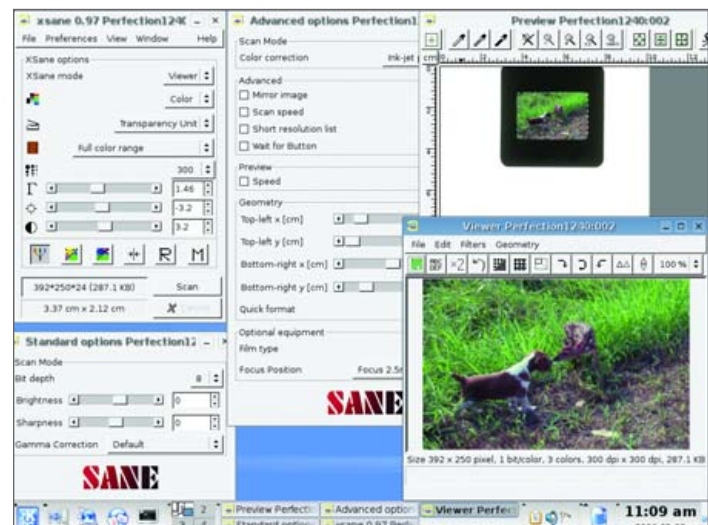


Figure 2. XSane with the Preview Image and Final Scan Image

Preview window. The preview scan displays in the Preview window, and you can use the mouse to select an area for the final scan. Clicking on the Scan button in the Program window initiates the scan with the resulting image displayed in an image viewer. The viewer has some limited abilities as far as editing the image, but its primary purpose is to proof the scan. Selecting the File menu item accesses the Save image menu item. From here, you can save the image in any one of a variety of formats.

A side note—for my scanner, in the Advanced Options section, there is the ability to focus on the glass or 2.5mm above the glass. For negatives in the the film holder or slides in thick mounts, the 2.5mm setting is best. You may want to experiment with this setting to see what works best.

XSane also works as a plugin to The GIMP. This means if you are using The GIMP for your image-editing chores, you can go to File→Acquire→XSane→Device Dialog from The GIMP and call up XSane directly. The final scanned image will be loaded directly into The GIMP for further editing.

digiKam

For working with pre-existing digital images, we look at two applications, digiKam for KDE users and F-Spot for GNOME users. Both serve as clients for the gPhoto2 libraries. gPhoto2 is an open-source project that has created a software interface for digital cameras. At the time of this writing, it supports approximately 600 cameras (see Resources). digiKam is found in the Graphics menu item on the KDE start menu. Clicking on the program opens a single window (Figure 3).

When you open digiKam the first time, you need to set up the camera. Make sure the camera is plugged in to the computer and turned on. Select Camera→Add Camera from the menu bar. First try Auto-Detect. If this does not work, click Add, and browse the list for your camera. Many cameras don't show up in the list but instead fall under the generic category of mass storage devices. Do not despair. When you hook up the camera to the computer and turn it on, it should show up on the desktop as a removable storage device. If this is the case and you cannot find the camera in the list, go to the right side of the window and select the port type (most likely USB) and camera mount path. The easiest way to find the mount path is to right-click the removable storage icon on the desktop and look under Properties.

Once you set up the camera, the fun begins. From the Camera menu item, select the camera or Directory Browse if it is seen as a storage device. This opens a window with the pictures stored on the camera. From here, you can use the buttons across the bottom to work with the images. When you download, you have the option of creating a new directory or storing to an existing directory. The pictures will show up in the right panel of the digiKam window. Opening an image loads it into an image viewer.

One stock capability of note found in the Fix menu item of the image viewer is Red Eye Reduction. The capabilities of the image viewer, as well as digiKam in general, can be extended by installing plugins. digiKam uses the KIPi (KDE Image Plugin Interface, see Resources) to manage plugins. The plugins include a module to work with RAW images. See the sidebar for information on RAW. It is possible to convert a RAW file to a JPEG, TIFF or PPM format. In versions 0.8.1 on, 8-bit RAW files can be opened in the image viewer, with 16-bit support promised in version 0.9.0.

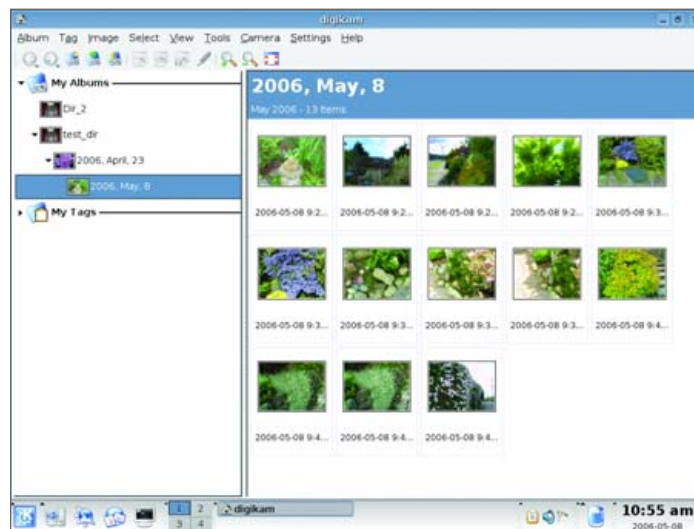


Figure 3. digiKam Showing Downloaded Images

F-Spot

F-Spot serves the same role in GNOME as digiKam does in KDE. It is not as far along in the development cycle as digiKam, so it is not quite as feature-rich. However, it does work well for pulling pictures from a camera and organizing them on your hard drive.

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

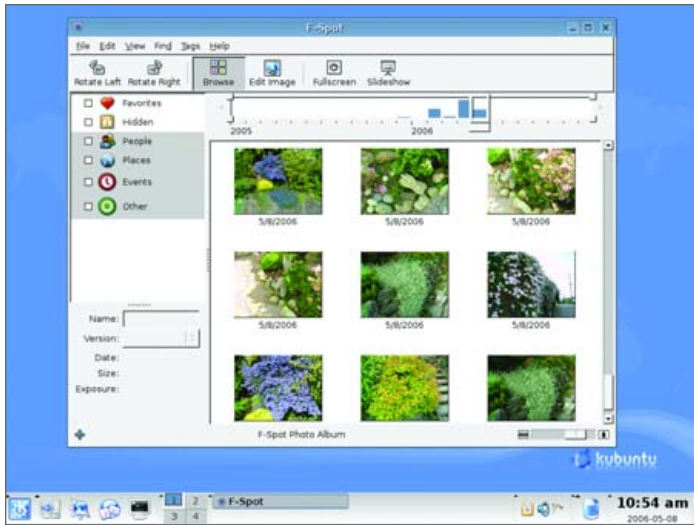


Figure 4. F-Spot showing downloaded images—note the time-line bar across the top.

Opening the program presents a single window with a couple panes (Figure 4). In the right pane are thumbnails of the pictures that have been downloaded. Right-clicking on an image brings up a menu that, among other things, allows you to tag the image. There is a set of predefined tags, and you can create your own tags using the menu bar item Tags. The tag labels are listed in the left pane, and selecting them filters the thumbnails accordingly.

An interesting feature of the thumbnail browser is the time-line bar across the top. It represents a visual indication of the number of pictures by month and year. Use the mouse to move the slider (as you do so, you will see a date appear above the slider) to the month and year you are interested in.

Image Editing—The GIMP

The discussion to this point has been concerned with getting pictures into the computer from an external source. Some mention has been made of manipulating images. Kooka, XSane, digiKam and F-Spot all can change the imported images to some degree,

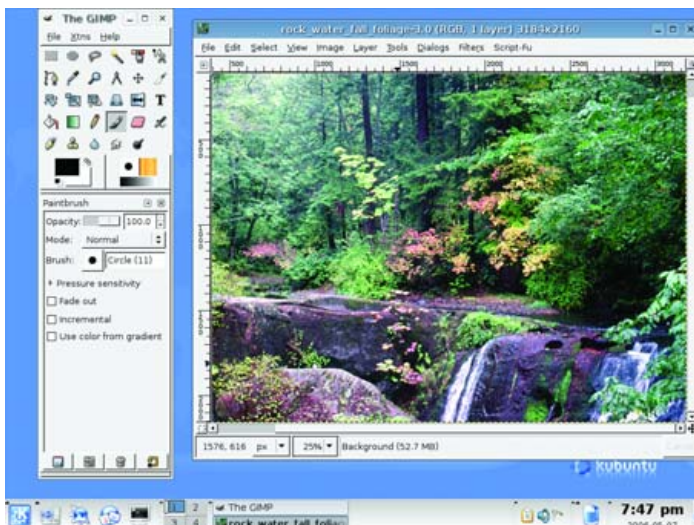


Figure 5. The GIMP with an Open Image

but to do more in-depth editing, it is necessary to use a dedicated image program. For this task, we look at The GIMP and CinePaint. The GIMP (GNU Image Manipulation Program) has been around since 1995 and is probably the most widely known open-source image program. CinePaint started as a fork of The GIMP in 1998, known as Film GIMP (also Hollywood GIMP and GIMP16) and acquired its current name in March 2003.

The GIMP follows the multiple window paradigm we saw earlier with XSane (Figure 5). The program window contains the toolbox. Placing the cursor over a tool icon brings up a help tip, and selecting an icon loads the tool options into the bottom half of the window. Once a tool is selected, moving the cursor over the image window shows the cursor pointer along with the tool icon, and left-clicking the mouse causes the tool action to take place. For example, selecting the eraser icon erases the area under the cursor according to the parameters selected in the tool options box (Figure 6).

Space does not permit a full exploration of the features in The GIMP, so I point out those I have found particularly useful. First, in an open image window, select Tools→Color Tools→Levels. A new window opens (Figure 7). Look for the Auto button, click on it and the image color levels are adjusted automatically (Figure 8). Often this is all you need to do to improve a picture. If the Auto correction does not work, next to the Auto button is a row of eye dropper icons representing the black point, grey point and white points of the image. Use one of the eye droppers to select a “known” color point in the picture, and the picture adjusts accordingly. You can undo your changes by clicking on Reset, so experiment away.

This leads me to another often-used feature of The GIMP: Ctrl-Z—the undo command. The extent to which you can undo is set in the Preferences item of the File menu. The two main parameters are the minimal number of undo levels and the maximum undo memory. The first specifies the number of undos to save, regardless of memory usage. The second causes undo to drop items once they exceed the memory limit, except those exempted by the first parameter. The undo command has some other limitations; see the manual for a full rundown, but it does allow you to experiment with relative impunity.

This is good place to insert a practical tip learned in the school of hard knocks. Do not work with your original image. The first

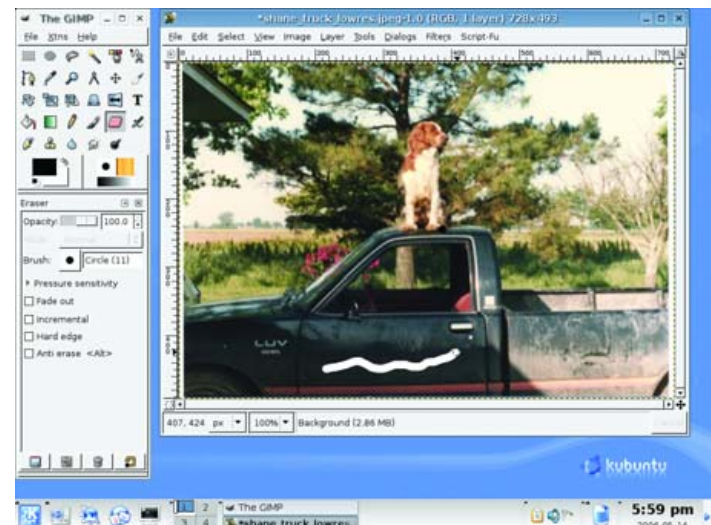


Figure 6. Using the Eraser

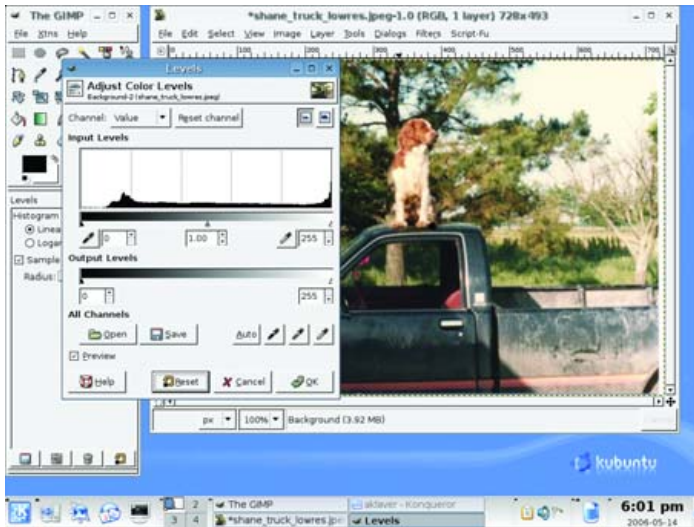


Figure 7. The GIMP—Auto-Adjust. Before

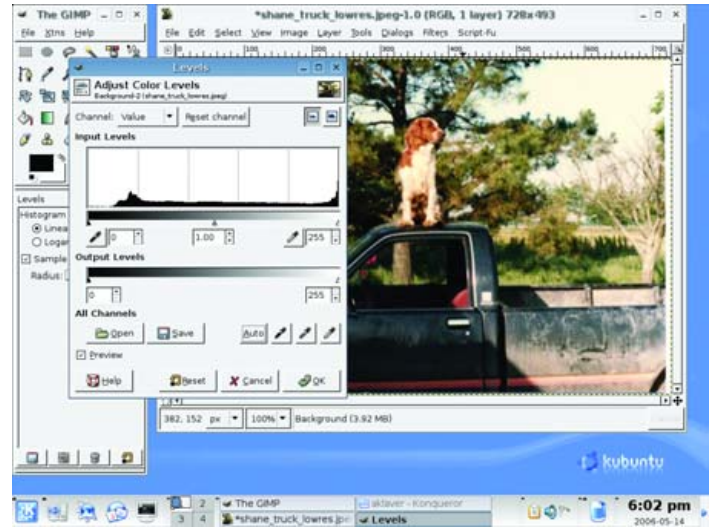


Figure 8. The GIMP—Auto-Adjust. After

step when working with an image is to create a copy and work with the copy. That way, when you paint yourself into a corner or commit an unrecoverable act, you still have the original to start over with.

The real flexibility comes through the plugin architecture of The GIMP. The program comes with a series of default plugins, and you can extend these by downloading others (see Resources). Plugins can show up in many areas, so the best way to identify what is available is to use the plugin browser located in the Xtns menu item of the program window. One optional plugin I have found particularly useful is UFRAW (see Resources). It allows you to get around the 8-bit limitation current versions of The GIMP have on RAW files. Basically, it loads a RAW file into its own viewer at its native

bitrate and allows you to adjust the image (Figure 9). After the image has been changed in the plugin, clicking OK mixes it down to the 8-bit format of The GIMP and opens a window for further work in The GIMP.

Versions of Software Discussed in This Article

- ▶ The GIMP 2.2.8
- ▶ CinePaint 0.19-1
- ▶ digiKam 0.7.2
- ▶ F-Spot 0.1.3
- ▶ Kooka 0.44
- ▶ XSane 0.97
- ▶ gimp-ufraw 0.5

DEDICATED SERVERS ROCK BOTTOM PRICING



\$59 /mo

64bit, 2.97GHz
Celeron D
Brand-Name Hardware

- » 1200GB/mo Bandwidth
- » 1GB RAM DDR400!
- » 160 GB SATA2 Hard Drive
- » FREE REBOOTS 24x7
- » FREE SUPPORT 24x7
- » 100Mbps Dedicated Port

ww.cari.net **carinet**
betterServers. betterService.

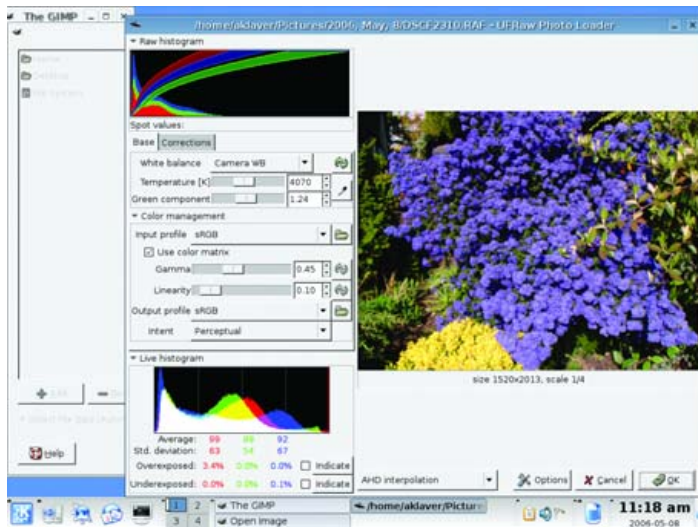


Figure 9. UFRaw Plugin for The GIMP

CinePaint

As mentioned previously, CinePaint is a fork of The GIMP. The program was developed to meet the needs of image retouchers in the movie industry. One of the requirements was the ability to

The Lowdown on RAW files

RAW images are, in essence, digital negatives. They represent the uncorrected output of the digital camera sensor. As such, they contain more photo information than a processed JPEG. For instance, a RAW will have bit depth of 12–14 per color versus that of 8 for a JPEG. The result is an image that allows for more editing opportunities. The downside is that this information increases the size of the files, so fewer images will fit on a memory card.

work with high-resolution images. Consequently, CinePaint has native support for RAW images and other high-bit formats. The trade-off is the inability to use GIMP plugins. There is work in progress to correct this, but it is in the development branch of the project. Still, the program shares many features with The GIMP, so if you are looking for some fairly advanced editing of high-bit images, it is a simple step up (Figure 10).

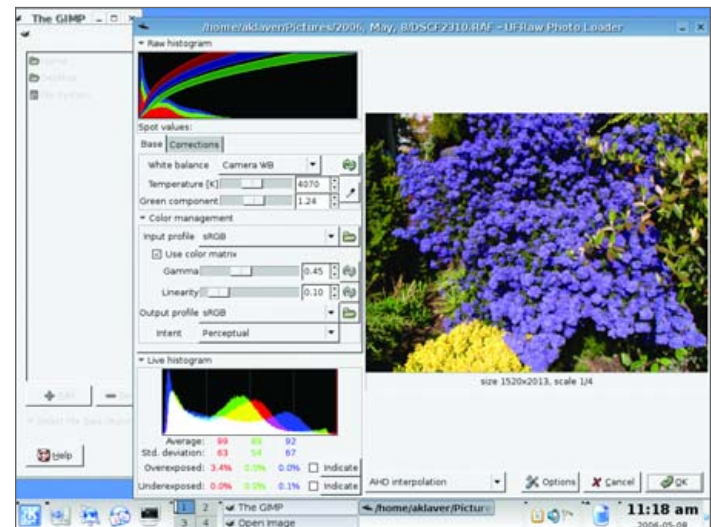


Figure 10. CinePaint with Same Image as Shown in Figure 9

Conclusion

It is possible for a relative beginner to work with digital images in Linux. The tools exist, and if not installed by default, are readily available. My personal work flow is to use XSane for scanning my print and slide collection, digiKam for bringing in images from my digital camera and The GIMP for retouching images. This article has touched on the basics of the programs explored. See the on-line Resources for a wealth of documentation for those who want to delve deeper. ■

Resources for this article: www.linuxjournal.com/article/9196.

Adrian Klaver is a Linux enthusiast with an interest in photography and a desire to make the former work with the latter.

LINUXWORLD[®]

CONFERENCE & EXPO

25-26 October 2006 • Olympia 2 • London

PLATINUM SPONSOR



Open. For Business.

EXPO

- See new products and services from over 60 exhibitors
- Meet developers and technical specialists in the .Org Village
- Take LPI examinations at a special visitor discount

Free-to-attend sessions include:

- Business Briefings
For Corporate and Senior IT managers looking to gain real life examples of advantages, disadvantages, applications and integration of Linux and Open Source in business.
- Great Linux Debate
- Seminars and Presentations

CONFERENCE

Featuring world-leading technical speakers discussing today's hottest topics. Aimed at technical and IT professionals the 2006 conference offers:

- 3 conference/master class streams
- Over 32 hours of content
- In-depth technical sessions
- Business focused topics
- Sessions from just £95

Book your **passport package** today at www.linuxworldexpo.co.uk/conference

LinuxWorld Conference & Expo is the premier event exclusively focused on Linux and open source solutions. As the world's most comprehensive marketplace for open source products and services, LinuxWorld provides business decision-makers and technical professionals with information and resources to implement Linux and open source solutions into business infrastructure and enterprise networks.

Register today for **FREE** queue jumping entry to the expo and save £15
www.linuxworldexpo.co.uk

EXPERT SPEAKERS INCLUDE:



Jon "maddog" Hall
Linux International



Ian Pratt
Kings' College Cambridge



Alan Cox
Red Hat



Bruce Perens
Sourcelabs



Jeremy Allison,
Novell



Chris DiBona
Google



Bill Weinberg
OSDL

Register now for updates at www.linuxworldexpo.co.uk

ORGANISED BY:

turret GROUP
International Business Publications & Exhibitions

LinuxWorld Conference & Expo is owned and licensed by IDG World Expo, which is a business unit of IDG, the world's leading technology media, research and event company. Copyright © 2006 IDG World Expo Corp. All rights reserved. LinuxWorld and LinuxWorld Conference & Expo are registered trademarks of International Data Group, Inc. E&OE. Programme may be subject to change. Correct at time of press.


IDG
WORLD EXPO

IDG
INTERNATIONAL DATA GROUP

MythTV Goes Amiss

Did MythTV get its name because the ability to install it easily is a myth?

NICHOLAS PETRELEY



Nick Petreley, Editor in Chief

I miss my TiVo. I had one of the first TiVo units ever made, and I kept it alive through various hardware failures and it still works today. I replaced the existing drive with a 160GB drive, which gives me virtually unlimited storage. Then my television failed, so I picked up a terrific bargain on an HDTV from www.newegg.com. The TiVo will work with this television, but it doesn't really do it justice.

So, I've decided to do something I've been meaning to do for years now. I'm building a MythTV box on Linux, which will serve as a multipurpose multimedia server and perhaps even take over the jobs of my aging all-purpose storage and IMAP e-mail server. I've put off this project for a number of reasons. I didn't want to invest in the extra hardware, and from everything I've read it is time consuming (read: royal pain) to get a MythTV box working.

All the testing I did for the Ultimate Linux Box issue left me with a lot of extra hardware to play with. So that problem was solved for me. The only thing I had to purchase was a

TV tuner card. I went for the Hauppauge WinTV-PVR-500MCE, which is basically two WinTV-PVR-150 tuners on a single card.

The WinTV-PVR-500 doesn't do HDTV, but there aren't many HDTV channels either. So the plan is to create a synergy between the MythTV box and my HDTV-capable cable box so I can get the best from both. In most cases, the MythTV box should change the channel on the cable box.

The first thing I did was test out the card in a PC I already had running, namely a box with an ASUS M2N32-SLI motherboard for the socket AM2 AMD64 processor. While I had the case open to install the Hauppauge card, I took the opportunity to install some new DDR2 memory modules, which I plan to test for speed improvements. After that, the PC wouldn't even boot to the setup screen.

I'd read that this motherboard is finicky about memory, thanks to BIOS issues, so I used the old memory to download and install the latest BIOS. That fixed the memory problem, but now every distribution of Linux I have crashed at boot time. I get the MP-BIOS "8254 timer not connected to IO-APIC" error that plagued some earlier Linux kernels. Although this may be a Linux kernel issue, it's more likely a problem with the new BIOS. I've never had a lot of luck with ASUS fixing BIOS bugs, so I hope the kernel developers find a way to work around the problem. Regardless, I added noapic to the boot command, and now Linux boots again.

By this time, I decide I don't want to mess with the ASUS box, so I build the computer with the slightly older, but still powerful hardware I planned to use for MythTV. The Knoppix distribution customized for MythTV called KnoppMyth is almost impossible to install, and it doesn't work when I get it installed. Kubuntu 6.06 installs cleanly, but I make the mistake of using a vanilla 2.6.17 kernel to add the Hauppauge drivers (it requires ivtv drivers from dl.ivtvdriver.org

and firmware from the Hauppauge site). Vanilla kernels break numerous things in Kubuntu, and the Hauppauge drivers simply won't compile properly anyway.

So I resign myself to using an existing Ubuntu/Kubuntu kernel. I found numerous how-to sites for getting the ivtv drivers working on Ubuntu 5, but nothing for Ubuntu/Kubuntu 6. Unfortunately, there are significant differences between these versions. The biggest difference is how the two versions handle third-party firmware. After experimenting for far too long, I find the shotgun approach solves my problem. I put copies of the driver modules and firmware files just about everywhere, including my sock drawer. Finally, the modules load and the /dev directory is now populated with the right devices.

I still have miles to go before I sleep, but the rest of the MythTV installation looks pretty simple, at least on paper (or should I say "on pixels"). However, the above experience hammered home something I've known about Linux and Linux distributions for a long time, but haven't spoken enough about lately. Almost every distribution has forked the Linux kernel and much more to deliver on its specific goals for a Linux distribution. I hesitate to call this a bad thing, even though it tends to rear its ugly head when you try to do something like install the drivers for MythTV, because updated how-to instructions almost always trail the release of new distributions by many months.

There are reasons why these distributions take such liberties. The basic foundation isn't good enough. Those are fighting words, I know. And I will leave you hanging with that flame bait. I have much more to say on this topic next month. ■

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.



Rackspace – Managed Hosting Backed by Fanatical Support™

Fast servers, secure data centers and maximum bandwidth are all well and good. In fact, we invest a lot of money in them every year. But we believe hosting enterprise class web sites and web applications takes more than technology. It takes Fanatical Support.

Fanatical Support isn't a clever slogan, but the day to day reality our customers experience working with us. It's how we have reimagined customer service to bring unprecedented responsiveness and value to everything we do for our customers. It starts the first time you talk with us. And it never ends.

Contact us to see how Fanatical Support works for you.

1.888.571.8976 or visit www.rackspace.com



Thanks for honoring us with the
2005 Linux Journal Readers' Choice Award for
"Favorite Web-Hosting Service"



The World's Fastest InfiniBand Switch

From a Company You've Trusted for 24 Years

Microway's **FasTree™** DDR InfiniBand switches run at 5GHz, twice as fast as the competition's SDR models. FasTree's non-blocking, flow-through architecture makes it possible to create 24 to 72 port modular fabrics which have lower latency than monolithic switches. They aggregate data modulo 24 instead of 12, improving nearest neighbor latency in fine grain problems and doubling the size of the largest three hop fat tree that can be built, from 288 to 576 ports. Larger fabrics can be created linking 576 port domains together.

Working with QLogic's InfiniPath InfiniBand Adapters, the number of hops required to move MPI messages between nodes is reduced, improving latency. The modular design makes them useful for SDR, DDR and future QDR InfiniBand fabrics, greatly extending their useful life. Please send email to fastree@microway.com to request our white paper entitled *Low Latency Modular Switches for InfiniBand*.



▲ 72 Port FasTree™ Configuration

Harness the power of 16 Opteron™ cores and 128 GB in 4U

Microway's **QuadPuter®** includes four or eight AMD dual core Opteron™ processors, 1350 Watt redundant power supply, and up to 8 redundant, hot swap hard drives—all in 4U. Dual core enables users to increase computing capacity without increasing power requirements, thereby providing the best performance per watt. Constructed with stainless steel, QuadPuter's RuggedRack™ architecture is designed to keep the processors and memory running cool and efficiently. Hard drives are cooled with external air and are front-mounted along with the power supply for easy access and removal. The RuggedRack™ with an 8-way motherboard, 8 drives, and up to 128 GB of memory is an excellent platform for power- and memory-hungry SMP applications.



Call us first at 508-746-7341 for quotes on clusters and storage solutions. Find testimonials and a list of satisfied customers at microway.com.

◀ QuadPuter® Navion™ with hot swap, redundant power and hard drives and four or eight dual core Opterons, offering the perfect balance between performance and density



508.746.7341 microway.com