

UBUNTU | ABIWORD | FIREFOX | RUBY | THUNDERBIRD | ASTERISK

# LINUX JOURNAL™

Since 1994: The Original Magazine of the Linux Community  
DECEMBER 2006 | ISSUE 152 | [www.linuxjournal.com](http://www.linuxjournal.com)

**Best of  
Seven Years**  
with Marcel Gagné



# Editors' Choice 2006



LINUX DISTRIBUTION • OFFICE SUITE • DESKTOP ENVIRONMENT • SPREADSHEET • WORD PROCESSOR • PRESENTATION SOFTWARE • WEB BROWSER • DATABASE • MAIL CLIENT • SECURITY TOOL • LANGUAGE • GAME/ENTERTAINMENT • WEB SERVER • WEB APPLICATION FRAMEWORK • COMMUNICATION TOOL • SOFTWARE DEVELOPMENT TOOL • BOOK • NOTEBOOK • GRAPHICS SOFTWARE • ADMIN SOFTWARE • HOSTING OR COLOCATION • MOBILE DEVICE • MORE

## PLUS

- LYX AND LULU SITTING IN A TREE
- TWO-FACTOR AUTHENTICATION CAN SAVE THE DAY
- HOW TO MEASURE MEANINGFUL LOAD AVERAGE
- SHED POUNDS OFF YOUR CLIENTS

AN  
**SSC**  
PUBLICATION

USA \$5.00  
CAN \$6.50

0 71486 03102 4 12

A standard 1D barcode with the numbers 0 71486 03102 4 12 printed below it.

# Enterprise and High-Performance Computing Under Your Control



## Industry Leading 2P and 4P x86 Computing

Innovative server technology with outstanding performance and memory scalability

### 2-way XtremeServer™ up to 64GB of memory



- AMD Opteron™ processors
- Featuring AMD Socket F(1207)
- Up to 64GB of DDR2 533/667 memory
- Up to 1.0TB SATA or 292GB SAS
- 1 PCI-X and 1 PCI-Express x16
- Hot-swappable drives
- ServerDome Management – IPMI 2.0
- Windows® or Linux OS

### 4-way XtremeServer™ up to 128GB of memory



- AMD Opteron™ processors
- Featuring AMD Socket F(1207)
- Up to 128GB of DDR2 533/667 memory
- Up to 3.0TB SATA or 876GB SAS
- 3 PCI-X and 2 PCI Express x16
- Redundant power supplies and fans
- Hot-swappable drives
- ServerDome Management – IPMI 2.0
- Windows® or Linux OS



**Keeping Your Cool in the Data Center while Consolidating & Virtualizing your IT Infrastructure**  
Go to [http://www.appro.com/whitepaper/White Papers.asp](http://www.appro.com/whitepaper/White%20Papers.asp)

AMD Opteron™ Processors - Quad-Core Ready - increase capacity without altering datacenter infrastructure  
- Best performance per-watt with energy-efficient DDR2 memory

For more information, please visit [www.appro.com](http://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](http://www.avocent.com/special) to download *Data Center Control: Guidelines to Achieve Centralized Management* white paper.



# CONTENTS

DECEMBER 2006  
Issue 152



## FEATURE

# 56 EDITORS' CHOICE 2006

*Linux Journal* editors pick their favorites. **LJ Staff**

### ON THE COVER

- Editor's Choice 2006, p. 56
- Best of Seven Years with Marcel Gagné, p. 32
- LyX and Lulu Sitting in a Tree, p. 64
- Two-Factor Authentication Can Save the Day, p. 70
- How to Measure Meaningful Load Average, p. 80
- Shed Pounds off Your Clients, p. 84

# 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](mailto:info@coyotepoint.com) for your free copy of the full Tolly Report.



**CoyotePoint**  
Systems Inc



# CONTENTS

## DECEMBER 2006

### Issue 152

#### COLUMNS

**24 REUVEN M. LERNER'S  
AT THE FORGE**  
Ajax Application Design

**32 MARCEL GAGNÉ'S  
COOKING WITH LINUX**  
Tonight's Menu: Diner's Choice!



**38 DAVE TAYLOR'S  
WORK THE SHELL**  
Breaking Numbers Down

**42 MICK BAUER'S  
PARANOID PENGUIN**  
Running Network Services under User-Mode  
Linux, Part II

**48 JON "MADDOG" HALL'S  
BEACHHEAD**  
Events for Suits

**50 DOC SEARLS'  
LINUX FOR SUITS**  
Greater Goods

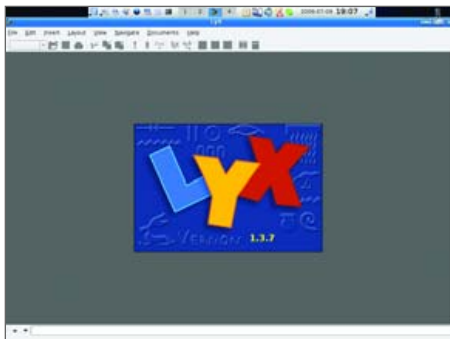
**96 NICHOLAS PETRELEY'S  
/VAR/OPINION**  
Plug This in Your Pipe and Smoke It

#### IN EVERY ISSUE

**12** LETTERS  
**16** UPFRONT  
**54** NEW PRODUCTS  
**81** ADVERTISERS INDEX

#### INDEPTH

**64 LYX AND LULU**  
LyX beats word processors for publishing.  
**Donald Emmack**



**70 TIGHTER SSH SECURITY WITH  
TWO-FACTOR AUTHENTICATION**  
Can a USB pendrive work as a smart card?  
**Paul Sery**

**74 A SERVER (ALMOST) OF YOUR OWN**  
Set up a virtual private server.  
**George Belotsky**

**80 EXAMINING LOAD AVERAGE**  
Can you trust CPU usage numbers?  
**Ray Walker**

**84 THINKING THIN**  
Do you really need a modern desktop PC  
to get your work done?  
**Lyle Frost**

**88 ACIDRIP—A GTK2 FRONT END  
TO MENCODER**  
DVDs are too fragile to keep without backups.  
**Daniel Bartholomew**

**92 FEDERATED DESKTOP AND FILE  
SERVER SEARCH WITH LIBFERRIS**  
Federate your indexes to enhance searches.  
**Ben Martin**



**84 THINKING THIN**

## Next Month

### EMBEDDED

We're in the dawn of the age of multicore processing. It's only a matter of time before multicore processors become ubiquitous in many types of embedded systems. How will this play out with respect to issues like real-time requirements? Be prepared by reading our journey into SMP and embedded systems requirements.

There's more. Want to create an unattended, encrypted, redundant network backup? We've got the solution. We also look at how to use Ajax to expand the semantic Web by adding a timeline component to your site.

**USPS** LINUX JOURNAL (ISSN 1075-3583) is published monthly by SSC Media Corp., 1752 NW Market Street #200, 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.



## Database + Backups + Archives

No matter how many peas are in your pod, Pogo Linux multi-tiered storage solutions can cover all of your needs. From ultra-fast primary storage, near line storage for disk to disk backup, to very affordable archival storage devices, Pogo Linux is your only source for high quality, high value hardware that speaks your languages. So climb out of your shell and give us a call today.

	+		+	
<b>PerformanceWare 1510T</b> <ul style="list-style-type: none"><li>• Compact and affordable</li><li>• Intel® Xeon® Dual-Core</li><li>• Contact us for price quote</li></ul>		<b>StorageWare SA350</b> <ul style="list-style-type: none"><li>• Up to 12TB in a single server!</li><li>• Intel® Xeon® Dual-Core</li><li>• Contact us for price quote</li></ul>		<b>StorageWare DA360</b> <ul style="list-style-type: none"><li>• Up to 16 SATA II Disk Drives</li><li>• Add this to those for capacities too large to mention here!</li><li>• Contact us for price quote</li></ul>




**Experience, Imagination, and Support.**  
**Pogo Linux is Hardware Solutions Built for Linux, Built for You.**

To get started, contact us at 888.828.POGO or [inquiries02@pogolinux.com](mailto:inquiries02@pogolinux.com)  
Pogo Linux, Inc. 701 Fifth Ave. Suite 6850, Seattle, WA 98104

[www.pogolinux.com](http://www.pogolinux.com)

Intel, Intel logo, Intel Inside logo, Pentium, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel corporation or its subsidiaries in the United States and other countries. For additional terms and conditions please visit [www.pogolinux.com](http://www.pogolinux.com)



# LINUX JOURNAL™

Since 1994: The Original Magazine of the Linux Community

## Digital Edition Now Available!

### Read it first

Get the latest issue before it  
hits the newsstand

### Keyword searchable

Find a topic or name  
in seconds

### Paperless archives


Download to your computer for  
convenient offline reading

### Same great magazine

Read each issue in  
high-quality PDF

### Try a Sample Issue!

[www.linuxjournal.com/digital](http://www.linuxjournal.com/digital)



# LINUX JOURNAL

### Editor in Chief

Nick Petreley, [ljeditor@ssc.com](mailto:ljeditor@ssc.com)

### Executive Editor

Jill Franklin  
[jjl@ssc.com](mailto:jjl@ssc.com)

### Senior Editor

Doc Searls  
[doc@ssc.com](mailto:doc@ssc.com)

### Art Director

Garrick Antikajian  
[garrick@ssc.com](mailto:garrick@ssc.com)

### Products Editor

James Gray  
[newproducts@ssc.com](mailto:newproducts@ssc.com)

### Editor Emeritus

Don Marti  
[dmarti@ssc.com](mailto:dmarti@ssc.com)

### Technical Editor

Michael Baxter  
[mab@cruzio.com](mailto:mab@cruzio.com)

### Senior Columnist

Reuven Lerner  
[reuven@lerner.co.il](mailto:reuven@lerner.co.il)

### Chef Français

Marcel Gagné  
[mggagne@salmar.com](mailto:mggagne@salmar.com)

### Security Editor

Mick Bauer  
[mick@visi.com](mailto: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 Kenney

### Proofreader

Geri Gale

### VP of Sales and Marketing

Carlie Fairchild  
[carlie@ssc.com](mailto:carlie@ssc.com)

### Marketing Manager

Rebecca Cassity  
[rebecca@ssc.com](mailto:rebecca@ssc.com)

### International Market Analyst

James Gray  
[jgray@ssc.com](mailto:jgray@ssc.com)

### Sales Coordinator

Lana Newlander  
[ads@ssc.com](mailto: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

### Advertising Inquiries [ads@ssc.com](mailto:ads@ssc.com)

### Publisher

Phil Hughes  
[phil@ssc.com](mailto:phil@ssc.com)

### Accountant

Candy Beauchamp  
[acct@ssc.com](mailto: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](mailto:subs@ssc.com)

URL: [www.linuxjournal.com](http://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., 1752 NW Market Street #200, 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.

**LINUX** is a registered trademark of Linus Torvalds.





Iraq

University of Baghdad



Kosovo

Pristina



Benin

Cotonou



Denmark

Djursland



Mongolia

Ulanbataar



Mozambique

Maputo



South Africa

Western Cape



Iraq

Sulaimaniyah



Italy

Apennine mountains



Czech Republic, The First MikroTik User Meeting, 154 participants from 37 countries

Prague

MikroTik around the world  
no matter what the conditions, mikrotik is routing your world. [www.mikrotik.com](http://www.mikrotik.com)

# 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**

### Upgrade Your System

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

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

- 150 Gigaflops computing power.
- Up to 128GB<sup>1</sup> 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.





# Monarch Computer.com™



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



Discover a powerful and exciting leap forward in your computing experience with AMD Athlon 64.



**WE ARE MAKING A DIFFERENCE. MONARCH HAS DONATED OVER \$50,000.00 YEAR TO DATE!**

**Help Fight Cancer, Monarch makes a \$1.00 donation for every Lance Armstrong bracelet. Simply select the Free Lance Armstrong bracelet when buying any AMD Athlon 64/FX/X2/Opteron 100/1000 series Combo, System or Barebones system.**

**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**

Asus M2NPV-VM w/ AMD Athlon™ 64 X2 3800+ (Dual-Core) (Socket AM2)

Only \$238

Asus M2N32-SLI Deluxe nF 590 w/AMD Athlon™ 64 X2 4600+ (Dual-Core) (Socket AM2)

Only \$438



Asus M2N32 WS Professional nF 590 w/AMD Athlon™ 64 FX-62 (Dual-Core) (Socket AM2)

Only \$1,047

Supermicro H8DAE-2 PCI-E/SATA/Reg-DDR2 w/2 x AMD Opteron™ 2210 (Dual-Core) & Heatsinks (Socket F)

Only \$1,399



Tyan S3992G3NR Thunder PCI-E/SATA/Reg-DDR2 w/2 x AMD Opteron™ 2216 (Dual-Core) & Heatsinks (Socket F)

Only \$2,028

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](http://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 DFI INFINITY NF ULTRAI-M2 AMD Athlon™ 64 3500+ (Socket AM2)

Starting @ \$345



Monarch Hornet Pro (Silver) w/270W PS Asus M2NPV-VM AMD Athlon™ processor 3200+ (Socket AM2)

Starting @ \$375

Antec Titan550 E-ATX Full-Tower/Server w/550W PS

Supermicro H8DAE-2 AMD Opteron™ 2212 (Dual-Core) (Socket F) & Thermaltake Heatsink/Fan 2 x 512 MB DDR2 (533) PC2-4200 REG ECC Buffalo Memory

Starting @ \$1,316

SuperMicro (Barebone) AS1021M-82B 1U 4-Way Dual-Core Opteron W/Hotswap SCSI & 1000W PS (Black) Supports up to 2 x AMD Opteron™ Dual-core 8000 series (Socket F) Processors

Starting @ \$1,999



**Lowest Prices on AMD!**

[www.monarchcomputer.com](http://www.monarchcomputer.com)

## Components and Upgrades

1000s of In-Stock PC Parts and Accessories

**AMD Athlon™ 64 X2 Dual Core Retail CPUs (939 and Socket AM2)**  
AMD Athlon™ 64 X2 4600+ (939) \$265.00  
AMD Athlon™ 64 X2 4800+ (939) \$329.99  
AMD Athlon™ 64 X2 3800+ (AM2) \$159.99  
AMD Athlon™ 64 X2 4200+ (AM2) \$188.99  
AMD Athlon™ 64 X2 4600+ (AM2) \$265.00  
AMD Athlon™ 64 X2 5200+ (AM2) \$479.00  
AMD Athlon™ 64 X2 3800+ 65W (AM2) \$200.00  
AMD Athlon™ 64 X2 4200+ 65W (AM2) \$230.00  
AMD Athlon™ 64 X2 4600+ 65W (AM2) \$296.00

**AMD Athlon™ 64 Single Core Retail Box CPUs (939 and Socket AM2)**  
AMD Athlon™ 64 3200+ (939) \$75.00  
AMD Athlon™ 64 3800+ (939) \$92.99  
AMD Athlon™ 64 3800+ (939) \$113.99  
AMD Athlon™ 64 4000+ (939) \$295.00  
AMD Athlon™ 64 3200+ (AM2) \$81.99  
AMD Athlon™ 64 3500+ (AM2) \$96.99  
AMD Athlon™ 64 3800+ (AM2) \$111.99

**Dual-Core AMD Opteron™ Processors OEM CPUs (939 pin)**  
AMD Opteron™ 165 1.8GHz \$199.00  
AMD Opteron™ 170 2.0GHz \$323.00  
AMD Opteron™ 175 2.2GHz \$426.00  
AMD Opteron™ 180 2.4GHz \$556.00  
AMD Opteron™ 185 2.6GHz \$696.00

**Next-Generation AMD Opteron™ Processors (2000 Series)**  
AMD Opteron™ Model 2210 \$300.00  
AMD Opteron™ Model 2212 \$415.00  
AMD Opteron™ Model 2214 \$572.00  
AMD Opteron™ Model 2216 \$730.00  
AMD Opteron™ Model 2218 \$882.00  
AMD Opteron™ Model 2220 SE \$1,210.00

**Next-Generation AMD Opteron™ Processors (8000 Series)**  
AMD Opteron™ Model 8212 \$869.00  
AMD Opteron™ Model 8214 \$1,192.00  
AMD Opteron™ Model 8216 \$1,497.00  
AMD Opteron™ Model 8218 \$2,169.00  
AMD Opteron™ Model 8220 SE \$2,589.00

**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 \$689.00  
AMD Opteron™ 256 3.0GHz \$855.00  
AMD Opteron™ 850 2.4GHz \$677.00




150439 Western Digital Caviar SE 320 GB SATA 8MB Cache 7200RPM (WD3200JD) \$97.99



150353 Western Digital Raptor X 150 GB SATA 16MB Cache 10K RPM (WD1500AHFD) \$239.99



100338 Enermax Liberty ELT500AWT ATX 2.2 500W SLI Power Supply \$96.99



101208 OCZ GameXStream 700W ATX 2.2/EPS/SLI Power Supply \$129.99



1501141 3Ware Escalade 9500S-8MI 8 Port SATA RAID Controller w/Battery Backup \$443.85



100329 Cooler Master Stacker STC-T01-UW1 E-ATX/BTX Full-Tower (Silver) \$139.99



141506 1 GB (2 pcs 512) DDR2 (667) PC2-5300 Patriot (PDC21G5300LLK) \$191.42



141801 512 MB DDR (400) PC-3200 Buffalo Select (DD4002-512/BR) \$74.88



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.

**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 [monarchcomputer.com](http://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 releases, 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 [monarchcomputer.com](http://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<sup>®</sup> 2 line of systems now features Next-Generation AMD Opteron<sup>™</sup> processors with DDR2 and AMD Virtualization<sup>™</sup>!**

**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<sup>™</sup> 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<sup>™</sup> 1000/2000/8000 Series processors**

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



LOOK FOR THIS SYMBOL  
 ON LINUX COMPATIBLE  
 MONARCH SYSTEMS!

The Official Linux Journal Web Server



**GET QUICK QUOTES**

Online or by phone:

[www.monarchcomputer.com](http://www.monarchcomputer.com)

**1-800-611-0875**

Paypal - Visa - Mastercard - Discover - AMEX

**Monarch Empro<sup>™</sup> 2**  
**4-Way 3U UNIWIDE 3546VA**  
**Custom Server**  
 Ask for Part #: 80488



Configuration Starting @  
**\$4,299!**

The AMD Opteron<sup>™</sup> 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.

**SELECTED COMPONENTS:**  
 Uniwide 3546VA 3U  
 Quad Opteron (Dual-Core Ready)  
 w/6 SATA Hotswap, Rails & 1000W RED. PS  
 Choose up to 4 AMD Opteron<sup>™</sup> 8000 Series  
 (Socket F) Processors  
 Choose up to 128GB of ECC/REG  
 DDR2 400/533/667 SDRAM Memory  
 Choice up to 6 SATA or SAS Hot-swap Drives

**The Ultimate Compact  
 8 Core / 4-Way Cluster**

Supports up to 4 (8000 Series)  
 Next-Generation AMD Opteron<sup>™</sup>  
 Processors

Up to 128 GB of DDR2 Memory

**Monarch Empro<sup>®</sup> 2**  
**Custom Workstation**



Ask for Part #: 80667

**AVAILABLE COMPONENTS:**  
 Choice of Tower Case  
 Choice of Tyan or SuperMicro Dual  
 Opteron (Socket F) Motherboard  
 Up to 2 AMD Opteron<sup>™</sup> 2000 Series  
 Processors (Dual-Core)  
 Up to 32 GB DDR2 (400) PC2-3200  
 Up to 6 Harddrives (SCSI or SATA)  
 Optional RAID Options  
 Optional Optical Drives w/Software  
 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

Custom configs  
 Starting @ **ONLY**  
**\$1,599!**



**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.



## Scanning for Hardware That Works

I enjoyed the October 2006 issue of *LJ*—especially the article “Digital Photography and Linux” by Adrian Klaver, as I was just starting to investigate what it would take to transfer my 1,000 or so 35mm slides onto a CD or DVD.

I have hit a snag right out of the box. When I tried to find scanners with manufacturer-supported drivers, the list was not long and the data was quite old. I did some additional searching and found a lot of messages railing about the support, or lack of, regarding various models. It occurs to me that a good subject for a future article would be how to determine if your “widget” is supported. I know the standard answers, but when the data is old and talks about items that are no longer available, it is not helpful. Regarding Adrian’s article, I know that we cannot expect Adrian, or most authors, to test on various hardware, but it certainly would be a big plus if *LJ* could provide an addendum to an author’s article that stated the process has been tried on the different hardware, or include a list of supported hardware that could be used to accomplish the tasks the article describes—not an exhaustive list but some representative items.

I know everything is a resource issue, and I want to congratulate you on an excellent magazine. The October 2006 issue hit two of my hot buttons right on time. Thank you, and keep up the good work.

—

**Jim**

## I’m a Good Driver

Several recent discussions about the Linux kernel have focused on the problems it has with drivers. The monolithic kernel makes drivers a part of the kernel, and it is becoming bloated. In addition, management of modules, where infrequently used drivers are often placed, is becoming problematic, and the concept of modules may be dropped. A recent article (A. S. Tannenbaum, J. N. Herder, H. Bos, “Can We Make Operating Systems Reliable and Secure”, *IEEE Computer*, May 2006) pointed out some advantages to placing the drivers in user space rather than in the kernel. Minix 3, which is closely related to Linux, does this with great success. Because the drivers are in user space, only the drivers that are needed by the system have to be loaded. Infrequently used drivers can be loaded on demand. If a driver fails or is co-opted by rogue software, it does not cause the system to fail, and it can be recovered simply by reloading a fresh copy.

This approach would seem to have some real advantages for Linux. I understand that some experimental work has been done. The kernel size would be reduced, modules would not be needed and reliability would be enhanced. In addition, vendors who currently do not supply drivers for Linux because of problems they perceive with the GPL vs. the proprietary nature of their products could issue proprietary drivers that would not in any way be subject to the GPL and would operate strictly as user applications, interfacing to the primitive

interrupt handlers and dispatches in the kernel. Developers would need to pay attention to some new security concerns and attack modes, but overall the approach may be more secure than the current one.

—

**Norman Worth**

## One Linux, Many Faces

I thought I’d relate something that happened to me recently. A friend of mine called me up, excitedly saying that he just got a new computer and was telling me all the amazing things it could do, running Windows, of course, and he was especially excited about being able to change the way it looked with something called themes.

So I told him to come over to my place, and when he did I said to him, “just watch what I can do”. So, I logged in to my Ubuntu with GNOME, then said, “watch this”, logged out and logged back in with KDE. Then, before he could say another word, I logged out and logged in with WindowMaker (which I personally use most often), saw his look of total confusion, and finally logged out and back in with Fluxbox. At this point, completely confused, he said to me, “Wait. You have all these different OSes on your machine?” I told him it was all just one Linux distro. He was astounded.

So, when we argue over which WM is better, remember, what is important is not what is better, but the fact that we have the choice to use what works best for each of us, which is, if you think about it, certainly a lot better than having a big corporation forcing proprietary software down our throats without the ability of choice. Just a little comment on the argument of what’s best. By the way, love the way the magazine looks and reads these days. You have my subscription for a long time to come!

—

**Jon Alexander**

## Optimal *awking*

I was reading the October 2006 *LJ*, and at a certain moment, during reading the article from Dave Taylor named “Analyzing Log Files”, I noticed some processor consuming order in his examples.

At a certain point he wrote the next command prompt to search for HTML files in the `access_log`:

```
awk '{ print }' access_log | sort | uniq -c \
| sort -rn | grep "\.html" | head
```

This command consumes at my system:

```
real    0m0.097s
user    0m0.084s
sys     0m0.020s
```

If you put the `grep` command right after the `awk` command, the filter consumes:

# Enterprise Open Source Solutions for Demanding Data Center Environments

- > Open Source Systems is one-stop-shopping for all your enterprise data center needs. From custom system design, to onsite project deployment teams, our expertise is at your immediate disposal – 24 hours a day.

One of the most desirable value added services in the industry is rapid system deployment. At Open Source Systems we use neither contract manufacturers nor system integrators, instead investing in a production facility rated for over a thousand machine builds per day. All systems are designed, built, configured, and tested in-house, which translates into quality unmatched in the industry, and the ability to deliver within days instead of weeks.

All solutions are built to your specifications, loaded with your choice of software, and backed by an industry leading 3 year warranty.

**Configurable Reliable Affordable**



**AMD**  
Solution Provider  
PLATINUM



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

## [ LETTERS ]

```
awk '{ print }' access_log | grep "\.html" | sort | uniq -c \
| sort -rn | head
```

```
real    0m0.042s
user    0m0.028s
sys     0m0.012s
```

The reason why this is faster can be explained that in the first filter, you will sort first the *whole* data set, and after that you remove with `grep` the non-.html entries. The second one (the one I suggest) removes first all the non-.html entries and will sort it afterwards.

In my daily life, I have to deal with IT forensics and data analysis. I have a lot of big data sets and prefer the fast-as-possible commands (order) to do my work. With this data set in the article it doesn't matter (it is a fraction of second faster), but with data sets of more than 1GB it does matter.

--  
**Pieter de Rijk**

**Dave Taylor replies:** *a great point. I do spend most of my time working with smaller data sets, but you're right that greps should always be as early as possible to cut down that data stream.*

### Singing the Unsung Sister

I just had the occasion to re-read Mr Petreley's "Separation of Church and Choice" (March 2006), and it couldn't be more timely, as you'll see. For the record, I began reading *LJ* long before I started subscribing to it. *LJ* is a very well-put magazine, but, for me, of very limited value. There's nothing wrong with your choice of having the magazine heavily slanted toward sysadmins, programmers and hackers. Alas! I don't fall, nor intend nor expect to, into any one of those categories. Even though eons ago I had my hands heavily into Cobol, that's water long passed under the bridge. Nowadays, my use of Linux is mainly in spreadsheets, some writing, e-mail and a growing interest in digital photography. In other words, I'm a "domestic" user.

I'm not going to suggest you change *LJ*'s direction; I may be many things, but not stupid. My suggestion would be for you to consider a *sister* publication on the level of (does it still exist?) *Smart Computing*. Something for "household" consumption—beginners and intermediates.

I wish I were not, but I am not, in the economic level to shell out \$50 without concern for something I come to regret. Hence, I can't help but dream *LJ* could be of more help to "pedestrian" Linux users like me. I imagine there are a lot more of "us" than "them"; though, no question, their monthly income may easily surpass my annual intake.

--  
**g.r.**

*We've had such a publication for some time! It's called TUX, and I'm proud to say I was the Editor in Chief of TUX before moving to Linux Journal. You can find it at [www.tuxmagazine.com](http://www.tuxmagazine.com).—Ed.*

### MythTV Arcanity

Your October 2006 *LJ* commentary was dead-on target. I've been attempting MythTV for more than a year, without success. I've tried two motherboards, many distros (Red Hat, Mandriva, Knoppix 4, KnoppMyth and the Debian/AMICUS project). I also subscribe to the MythTV mail group. But, something, somewhere, always fails during the build. I freely admit that my own ineptitude plays a significant part, but I agree that building Myth is much more difficult than it should be. After all, the ATI AIW cards have been performing similar tasks since the early days of Win98 SE.

I built a box based on an Athlon64 board, with an NVIDIA graphics card, SB Audigy2 ZS sound, a Hauppauge PVR250 and a pcHDTV card, plus three drives with 800GB total space. All it needs is a functioning system. Maybe someday....

--  
**Joe O. Marcom**

*I finally got MythTV working fairly well, but the quality is limited by the available tuner and capture cards. There is only one hope that it will ever work as well as the built-in PVR in my HDTV cable box. The as-yet unreleased HDTV cable cards promise to capture HDTV just like a cable box. Let's hope they work, and that there will be Linux drivers for them.—Ed.*

### A Savage Take on Savage 2

I am writing to express my great disappointment in your article on *Savage 2* (September 2006). It seems that no one bothered researching S2 Games and its previous business practice before publishing this free advertisement for it.

When *Savage 1* came out years ago, I rushed to buy it because S2 Games supported Linux. The game installed and ran on Linux. Life was good...for two or three months. Then when a required patch came out, S2 Games never released a Linux version. *Savage*, being an on-line-only game, requires the same version. This kept all Linux users out in the cold with a worthless non-playable game. Myself and hundreds of others posted on the support forums, e-mailed tech support—all to no avail. S2 Games dropped Linux support and didn't even bother to respond to users. If you do a quick search, you will see the big stink over this. And, I'm sure someone will reply with "but it can work through a third-party hack" comment. If that's what you want to call supporting Linux, feel free to buy *Savage 2*. My vote, as always, will be with my money, and it won't go to S2 Games.

I am very disappointed that *Linux Journal* would publish such good things about S2 Games without doing research first into the company and its previous actions.

--  
**Greg**

*Sorry to fulfill your prediction, but it can run with a third-party hack. The S2 Games site credits Evolved Clan Community for continuing the support for Savage 1 on Linux and provides the appropriate link.—Ed.*

# LINUX JOURNAL

## At Your Service

### MAGAZINE

**PRINT SUBSCRIPTIONS:** Renewing your subscription, changing your address, paying your invoice, viewing your account details or other subscription inquiries can instantly be done on-line, [www.linuxjournal.com/subs](http://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](mailto: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.

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

**LETTERS TO THE EDITOR:** We welcome your letters and encourage you to submit them to [ljeditor@ssc.com](mailto: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](http://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](http://www.linuxjournal.com/advertising). Contact us directly for further information, [ads@linuxjournal.com](mailto: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](http://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](mailto: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](http://www.linuxjournal.com). Subscribe for free today, [www.linuxjournal.com/enewsletters](http://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](http://www.EmperorLinux.com) or call 1-888-651-6686 for details.

## EmperorLinux

...where Linux & laptops converge

[www.EmperorLinux.com/1-888-651-6686](http://www.EmperorLinux.com/1-888-651-6686)

Model specifications and availability may vary.



## diff -u

WHAT'S NEW  
IN KERNEL  
DEVELOPMENT

Alan Cox, Jeff Garzik and others have unveiled a plan to do away with the IDE subsystem entirely and completely replace it with libata. This will not happen immediately, and it will not happen all at once, but the plan does seem to have universal support. Even the creator of the original IDE subsystem, Mark Lord, thinks this is the way to go. For the immediate future, all that's happening is that more code will

merge from Andrew Morton's -mm tree into the official kernel, and users will have the option to use that improved support for various hardware if they so choose. The ultimate removal of the IDE subsystem is undoubtedly years away. Alan's recent announcement is only one step down a long road.

The fork from ext3 to ext4 is a reality. Once upon a time, adding features like journaling to ext2 was considered so invasive that folks had to fork ext3 in order to continue that kind of development. Now, the ext3 developers have had to take the same steps in order to add invasive features, such as extents and large block sizes, to the code. Linus Torvalds has stood firm on the idea that the most relied-upon filesystems should not in themselves undergo significant development, but should be rock solid and totally dependable. Now, the new ext4 code is on the fast track to being included in the official kernel. Whether it will ever be as popular as ext3 remains to be seen. Meanwhile, folks like Hans Reiser feel that ext4's easy entrance into the official tree is just further proof of the favoritism he feels is practiced within kernel development. What he doesn't understand is that intelligence and coding ability are only part of the kernel development culture. After all, one of Linus' great discoveries was that everyone has the ability to contribute, and kernel developers don't all have to be uber-hackers. They don't even have to be particularly nice guys, as Alexander Viro and others proudly proclaim. But, they do have to respond to feedback and present their work in more or less standard ways. The more they can be trusted to "do the right thing", the easier it is to get their code into the kernel.

Adrian Bunk will maintain the 2.6.16 kernel as a new stable tree. This has roughly the same appearance as if we still had the old even/odd stability model, Linus had forked 2.7 for intensive development, and Adrian were going to maintain 2.6 for stability. The only difference, it seems, are the names of the trees and the fact that Linus will not be stabilizing the stable tree himself for any length of time. Adrian's work on 2.6.16 will hopefully solve some of the issues users have had with the w.x.y.z stable tree maintained by Greg Kroah-Hartman and Chris Wright. That tree, although aiming for run-time stability, did nothing to prevent interface changes between 2.6 versions. Interface stability is not addressed at all by that effort, while the 2.6.16 interfaces will not change under Adrian's maintainership.

Pavel Machek has released a driver for ThinkPad fingerprint sensors. So far, users have reported good success with it, though at the moment, it does seem to have some easy-to-trigger failure modes. The big question for Pavel is whether to leave this as a user-space tool or to migrate it into the kernel proper. This is an interesting case, because typically anything that can reasonably be left outside the kernel, would be. Although at the same time, it is also typical to keep hardware support inside the kernel, with few exceptions. The direction of Pavel's code may influence where other drivers will live in the future as well.

Keith Packard from Intel has announced open-source drivers for Intel 965 Express Chipset family graphics controllers, as part of ongoing work by the Intel Open Source Technology Center. Intel seems to be doing the right thing here, acknowledging that the code needs testing and bug fixes, and inviting kernel folks to participate in development. One interesting detail quick to be noticed on the kernel mailing list is that the code seems to be written to interface with an unavailable binary blob, intel\_hal.so, if available. Keith explained, "This module contains stuff that Intel can't publish in source form, like Macrovision register stuff and other trade secrets. It's optional, so if you don't want to use a binary module, you don't get to use code written by Intel agents for these features....The driver remains completely functional in the absence of the binary piece and, in fact, has no reduction in functionality from previous driver releases."

—ZACK BROWN

Microsoft's  
New Promise

In the Free Software and Open Source worlds, licensing has always been a big deal. Choice of license has a direct effect on the usefulness of code bases, and on their market growth as well.

Some code, however, makes use of standards that are open, yet to some degree, proprietary. Those degrees are often controlled by patents. Lately, much lawyerly thinking has gone into making those standards useful to development efforts and to disarming the patents involved. One of these—perhaps the first—is the Microsoft Open Specification Promise. The Promise is short on legalese, yet too long to describe here, beyond saying it's about what Microsoft won't sue others for, providing others don't sue Microsoft. Lawrence Rosen, author of *Open Source Licensing: Software Freedom and Intellectual Property Law* (Prentice Hall 2004), says the Promise "...enables the Open Source community to implement these standard specifications without having to pay any royalties to Microsoft or sign a license agreement. I'm pleased that this OSP is compatible with free and open-source licenses."

The first standards in question involve SOAP and a variety of protocols from the WS-\* portfolio. At the time of the Promise's announcement, in mid-September, approving public statements were made by a variety of folks on the open-source side of the table. These include Mark Webbink, Deputy General Counsel of Red Hat, and R.L. "Bob" Morgan, Senior Technology Architect at the University of Washington.

Before bringing the defenses up, bear in mind that this promise has been hammered out through collaboration between open-friendly folks inside Microsoft and countless cooperative conversations with folks from Red Hat, Mozilla/Firefox, XRI/XDI, OpenID, LID, Sxip, Higgins, VeriSign and others, including customer-side entities such as North Carolina State University. The conversation has a name: OSIS, for Open Source Identity Selector (or something like that...even the initialism is open to change). Everybody involved is interested in developing open-source implementations—or products that interoperate with—Microsoft's CardSpace, an identity selector that will be released with Vista and which Kim Cameron and others at Microsoft have for several years been trying to make as interoperable as possible in a world where many other identity systems will be in use. (See the Identity Metasystem article from the September 2005 issue of *Linux Journal*.) But, the Promise may end up extending to other standards in other areas as well.

Disclosure: I've been involved in these discussions and have worked for some time to help make them happen and to move forward. It's clear we are at the beginning of something here, not an end.

Of course, the whole matter is open for input, debate and adjustment. To help with that, here are some links:


The Promise: [www.microsoft.com/interop/osp](http://www.microsoft.com/interop/osp)

OSIS: [osis.netmesh.org](http://osis.netmesh.org)

Kim Cameron: [identityblog.com](http://identityblog.com)

Let us know what you think.

—DOC SEARLS



# What could YOU do with **216** 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, December 2006

1. Millions of students in the state of Kerala, India: **1.5**
2. Number of government and government-aided schools in Kerala: **2,650**
3. Percentage of Kerala schools that will use or switch to free software on Linux: **100**
4. Thousands of Kerala high-school teachers being trained on Linux: **56**
5. Weeks between Richard M. Stallman's visit to Kerala and the state's decision to switch to free software and GNU/Linux: **2**
6. Percentage growth rate of the Indian economy: **8**
7. Percentage growth rate of the Chinese economy: **>10**
8. Billions of dollars Intel plans to invest over the next five years in its "World Ahead" program for emerging markets: **1**
9. Percentage of the world's population AMD wants to see connected to the Internet by 2015: **50**
10. Billions of people with annual incomes less than \$4,000 per year: **3.8**
11. Billions of people with annual incomes of \$4,000–\$20,000 per year: **1.5**
12. Percentage yearly growth of Cisco's network load: **100**
13. Projected percentage yearly growth of Cisco's network load: **300–500**
14. Cisco annual percentage growth rate in emerging markets: **30**
15. Percentage of new Cisco employees being hired in emerging markets: **12**
16. Percentage of television programs Cisco expects will be broadcast over the Net in the future: **100**
17. Current percentage contribution of emerging markets to Cisco's revenues: **10**
18. Projected maximum future percentage contribution of emerging markets to Cisco's revenues: **40**
19. Millions of dollars Google has donated to the One Laptop Per Child Project: **2**
20. Number of cities in Africa that Google intends to connect fully with a wireless network: **7**

**Sources:** 1–5: rediff news | 6, 7: *New York Times* | 8–11: *Electronic Engineering Times* | 12–18: World Resources Institute, reporting on a speech by Cisco CEO John Chambers | 19, 20: icicemac.com

—Doc Searls

## Another Open Letter to Bill Gates

Dear Bill,

I hope you are enjoying yourself in retirement. I (of course) am busy as ever trying to promote free software.

I read recently in the *Wall Street Journal* that you too have discovered the freedom of information sharing! I read in an article that you are insisting that researchers who receive your funding share their data, tools and results with each other. Awesome! I know that you may think this is another one of those "innovations" that you have come up with, but I have to tell you that this is the very core of the Free Software movement, and it has been going on for more than 35 years.

I remember back in 1969 when I was a student at Drexel University. I found some computers in the basement of Drexel's main buildings that did not come with software. In order to use these computers, I either had to write the software or buy it.

A single copy of a compiler for some language might cost \$100,000 US in those days, and that was when a hundred thousand dollars was a lot of money! I could not afford that on my small stipend for food and beer. But there were people in the Digital Equipment User's Society that wrote software and contributed it to the society's library for distribution to other people. It was the study of this software that allowed me to move into computer science. I have never forgotten that.

Of course, you may not have had the same enlightening experience. You went to Harvard and probably could afford the compilers of those days—or maybe you just used other people's machines and compilers to do your work.


In any case, as I left college and went out into the real world, I knew that working as a team is better than working alone, so I continued to push sharing code segments and even whole programs in order to make the industry move forward.

I just got a couple of great ideas!

1. All of the software you fund should be Free Software.
2. Use only Free Software in your own work.
3. Buy medical equipment only if it is supported by Free Software.

I am sure you will see how these fit into the basis of your new endeavors.

Warmest regards,  
Jon "maddog" Hall



# Don't browse the Web

# Create it

More **productive.**  
**Cooler** than you  
remember.

# Try Opera.

Download the  
**FREE** Opera browser  
[www.opera.com](http://www.opera.com)

*Opera 9 - your web, your choice*



# Soweto: Power from the People

In the spring of 2005, I attended LinuxWorld in Johannesburg, South Africa. It was not my first time to South Africa, but this time, instead of going to game reserves, I took a different trip after the event. A gentleman I had met at LinuxWorld introduced me to Soweto.

Soweto is a township outside of Johannesburg. Before apartheid ended, it was a township mostly of very poor black people. On June 16, 1976, students were killed in riots in Soweto that led to the beginning of the end for that rule of government called "apartheid".

My guide and I went on a little tour, first to the Photography Museum/School of Alf Kumalo. Dr Kumalo, who often risked his life to get photographs that illustrated for the world the issues of apartheid, was now using his talents and resources to teach young people in Soweto how to be photographers. I saw they were using Adobe's Photoshop to manipulate the digital images that the students were taking for composition training. I pointed out that they should use GIMP instead, because the students were unlikely to be able to afford Adobe Photoshop at home, and therefore they would have to pirate Photoshop. I promised Paballo Thekiso, the tutor that was my guide to the museum and school, that I would send them a copy of GIMP and a book on how to use GIMP from the USA.

Then we left and went to the Soweto museum. I read a lot of the information about what happened there, and as we looked out across the natural bowl-shaped valley, I mentioned to my guide (who was increasingly becoming my friend) that this would be a wonderful place for a mesh network to deliver Internet services to the entire township. We also talked about the benefits of FOSS and how there were no limitations to what students could learn, other than their own desires, assuming they had access to computer equipment and the Internet. I told him of several "success stories" for this concept, including one about a person who had been programming the kernel since the age of 12 and one about a person who had put out his own distribution at the age of 14. I bet him that there were people in Soweto who could "do Linux", given the opportunity.

We went to the house where his mother grew up and where his nephew still lived. It was a two-room house, and he talked about how there were sometimes two or three generations living in the same house. Although he was glad his nephew was doing well in school and sports, he was afraid that the nephew might turn to drugs, given the environment that still prevailed in Soweto.

We finally had dinner in a great outdoor restaurant on the edge of the township, where I experienced some of the local food and entrepreneurship that was happening there.

Then we drove back to Johannesburg, and I flew home. What I did not know was that my guide (and now friend) was Nhlanhla Mabaso, Open Source Center Manager for the Meraka Institute, and that he had been listening to me.

During the next year, I tried to send two books and two open CDs to the museum on two differ-

ent occasions, and neither time did they get through. Eventually, toward the end of the year I got them back, with a note on them saying that they were "undeliverable". It cost me more than \$150 US to buy the books and mail them. When I got them back, I was angry. Undaunted, I built into my 2006 LinuxWorld Johannesburg schedule some time to travel to the Photography Museum and to carry the books and CDs with me.

When I got my schedule from Aldean Prior, Director of Exhibits for Africa, the LinuxWorld producer, I noticed there was built in to it a trip to the Satellite Open Source Research Center that was opening as part of the Meraka Institute. I did not think anything about it, but I did keep asking to "go back to Soweto" so I could deliver the books to the museum. When I got to Johannesburg, I found out that the Satellite Center was in Soweto, and that I was invited to speak at the opening.

Nhlanhla drove me to Soweto that morning. We visited the museum, gave the books and CDs to Paballo and found out that he knew one of the people in the center. We invited Paballo to the first training class the next day, so he could establish contacts and learn more about Free Software.

The President and CEO of the Council for Scientific and Industrial Research (CSIR), Dr Sibusiso Sibisi, spoke during the opening, as well as Dr Ntsika Msimang and several other dignitaries. Dr Msimang is managing the new center. After the rest of the ceremonies were over, Dr Sibisi came over to me and said, "You have no idea how influential your words were."

Apparently, my conversation the previous year had inspired Nhlanhla to go back to the Meraka Institute and make a presentation for investigating the potential of opening an open-source research center and training facility in Soweto and also to start to set up a mesh network for the township. His team went to Soweto in August of that year and discovered (during a presentation on FOSS) that several of the young people in the audience already knew about Linux, could work with it and that one young man named Bongani Hlope was doing kernel programming as a hobby and conversing with Linus Torvalds via e-mail on kernel issues. Another person named Kgabo Sepuru was running a FOSS consulting service out of his house in Soweto.

In addition, Nhlanhla's team found that local people already had started to set up a mesh network in their broadband-deprived area.

On my second day at the satellite center, they had a day of training. That day started off a bit slowly because it was their first day, but I believe things will get better as they get settled in. Dr Msimang seems to be a competent, enthusiastic director, and I contributed the first book to their library.

In the meantime, Paballo, the tutor at the museum with whom I had been corresponding via e-mail all this time, got really excited about the rest of FOSS and said that he was going to learn Linux and teach it to the rest of the photography students. I promised to send some more books on other aspects of digital photography and image

rendering using free and open-source software.

I have had about five or six people come up to me in my life and say, "I listened to you, followed your advice on FOSS, started my own company and now I am a millionaire" or "I listened to you and it changed my life", but this was the first time I actually have seen direct action to this extent on something I said almost in passing. I can't take credit (nor do I want to) for the hard work that Nhlanhla and the rest of the staff put into making the center a reality, but it sure felt good to have someone like Dr Sibisi tell me those words.

Each of us affects the people around us with our every thought and deed. I often tell people that if they want to see the most influential person in free software, just look in the mirror when they get up. Lots of people do not believe what I say, and others do. Sometimes the effect of what we do and say goes way beyond what we know. I was fortunate enough to see how influential my words were, and therefore, I encourage others also to speak out and experience the same thing.

—JON "MADDOG" HALL

## They Said It

**When brokers turn into toll-takers, it's time to throw the bums out.**

—Britt Blaser, from a conversation

**Make no decision out of fear.**

—Bruce Sterling, from a speech at SXSW 2006

**We have decided that we will use only free software for computer education in Kerala schools. We have implemented the Linux platform in high schools; it will be implemented in other schools step by step....Our policy is to migrate computer education to free software platforms. We want to make Kerala the FOSS (Free and Open Source Software) destination in India. That is all.**

—M A Baby, Education Minister, Kerala, India, [www.rediff.com/money/2006/sep/02microsoft.htm](http://www.rediff.com/money/2006/sep/02microsoft.htm)

**We are getting lots of enquiries and orders for pre-loaded Linux operating systems. The hardware sales have gone up because of this.**

—P K Hari Krishnan, President, Kerala Computer Manufacturers' and Dealers' Association

**If there isn't enough food in the fridge, do you say "the store must be down"?**

—Greg Elin, at a conference

**What we need is an open-source, open hardware, wireless implementation for unlicensed spectrum. It can be done. And if it is, it will blow WiMAX out of the water and change the world.**

—Thomas A. Freeburg, at a conference

# 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](http://www.mbx.com) or call 1-800-681-0016 today.



[www.mbx.com](http://www.mbx.com) | 1-800-681-0016

**NEW!**

Tiny WiFi Controller  
boots Linux in 1.1 seconds




**\$119**  
CPU board only

**\$249**  
as shown


quantity  
discounts  
start at  
10 units

 **200 MHz CPU**

- TS-7400 CPU board
- Low power, low heat, long life
- Up to 128MB on-board Flash
- Up to 128M SDRAM
- SD Flash Card socket
- 1 external USB port
- 1 10/100 Ethernet
- 802.11g internal WiFi option
- One piece, rugged aluminum enclosure option measures 1.1" x 4.9" x 3.1"

 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

## [ UPFRONT ]



# PROGRESS TOWARD THE HACKABLE KIDTOP

One Laptop Per Child ([laptop.org](http://laptop.org)) hit the news in January 2005 at Davos, when Nicholas Negroponte of the MIT Media Lab announced plans for a "\$100 laptop" in a quantity of 100 million, to "revolutionize the way we educate the world's children". Formal plans for the project were announced in August, and we covered it for the first time in the November 2005 issue of *Linux Journal*.

Since then, much progress has been made. Jim Gettys ([www.handhelds.org/People/jg.html](http://www.handhelds.org/People/jg.html))—prime mover behind the X Window System, handhelds.org ([www.handhelds.org](http://www.handhelds.org)) and earlier fun projects like the Unobtainium (a wild hack on Compaq's original iPAQ)—is now VP of Software Engineering. And, there are prototypes. The current generation is fitted in bright orange and green and features rectangular bunny ears (802.11s mesh network antennae); a lid that twists and flips to form a pad; a dual-mode display that the project wiki says, "can readily be mass produced in standard LCD factories, with no process changes" and that "has higher resolution than 95% of the laptop displays on the market today, approximately 1/7th the power consumption, 1/3rd the price,

sunlight readability and room-light readability with the backlight off".

Although the units are designed for kids (the keyboard is 6/10 the size of an adult one), they're also made to hack. Writing on his blog ([www.ethanzuckerman.com/blog/?p=824](http://www.ethanzuckerman.com/blog/?p=824)), Ethan Zuckerman says:

...the 500 prototype boards currently built come with a VGA jack soldered on, but production models will leave the jack leads etched on the board, though unpopulated. Want to turn a laptop into a device that can drive an external monitor? Solder one on. Also on the board, but unpopulated, will be connectors for additional RAM and Flash memory, as well as a mini-PCI slot.

And, although some large companies (AMD, Google, News Corp., and Red Hat) are involved, the whole project is very much a work in progress, and it's open to interest and help with what promises to become the most widespread and good-hearted Linux deployment on earth.

—DOC SEARLS



We use our stuff.

visit our TS-7200 powered website at

[www.embeddedARM.com/7400wifi](http://www.embeddedARM.com/7400wifi)

(480) 837-5200



# AMD® Dual Core to 16 Processors!



Left is the Xi® MTower™ 2P64 workstation equipped with NVIDIA® nForce®5 Professional Chipset with HyperTransport™, dual PCIe™ & SLI™ capable.. Above is the Xi BladeRAIDER™ 64, 1U blade cluster, 2/4way Scalable, for render farm and parallel computing. Right is the Xi NetRAIDER™ 64, classic pedestal server with up to 6 Terabyte of RAID, hot swap cached disk space. All are equipped with the 64Bit AMD Athlon & Opteron™ Processors.

## Powerful & Quiet w/DDR2 800!

@Xi® Computer we lead the second generation of 64Bit Workstations, Servers and Clusters based on the 64Bit AMD Athlon™ and Opteron™ Processors. 32-Bit Compatible, 64-bit ready. Dual Core and soon Quad Core, with the fastest NVIDIA® Quadro & SLI™ and ATI® Fire-GL for maximum performance CAD/CAE/VIDEO Editing and Simulations. On the multiprocessor workstation & server front we now offer up to 8x64bit X2 CPU, 64GB of 16 channels ECC RAM and over 6 Tera-Bytes of RAID Data Storage to satisfy the most demanding computational models and the most pressing and critical real time requirements of our Scientific, HPC and DCC users. The Power Choice is @Xi Computer!



### Xi MTower™ 64 Xi MTower™ 2P64 NetRAIDER™ 64

- |                                 |                                   |                                  |
|---------------------------------|-----------------------------------|----------------------------------|
| - AMD Athlon™ 64 x2 Processor   | - 2xAMD Opteron™ x2 Processors    | - 2/4xAMD Opteron™ x2 Processors |
| - 2xPCI-E Nvidia® nForce®5 SLI™ | - 2xPCI-E Nvidia® nForce®Pro SLI™ | - 2xPCI-E Nvidia® nForce®Pro™    |
| - 1GB DDR2 800MHz to 4GB        | - 2GB DDR2 667 ECC to 16GB        | - 4GB DDR2 667 ECC to 16/32GB    |
| - NVIDIA EVGA® 7300GS 256MB     | - NVIDIA Quadro® Fx1500 PCI-E     | - 8 Ch. SATA2 CACHE RAID Ctr.    |
| - 160GB 7.2K SATA2 8MB cache    | - 160GB 7.2K SATA2 8MB cache      | - One TeraB Hswap RAID Mode5     |
| - 19" LCD 1280x1024NR w/Spkr.   | - 20.1" Viewsonic® VX2025 1680NR  | - 2x 1GB Ethernet (10/100/1000)  |
| - Xi MTower 600wPS 16xDVD-DL    | - Xi MTower 600wPS 16xDVD-DL      | - Xi 21Bays or 4U w/Dual HS PS   |
| - MS Window® XP or LINUX® RH    | - MS Window® XP-Pro or Linux®     | - RedHat® Linux®/MS® XP-PRO64    |

Xi MTower 64 x2 4200+ \$1,499

Xi MTower 2P64 2x2212 \$2,749

Xi NetRAIDER 64 2x2212 \$ 5,499

Xi MTower 64 x2 FX-62 \$2,249

Xi MTower 2P64 2x2220 \$5,299

Xi NetRAIDER 64 4x8220 \$16,999



### Dual & Quad monitor video upgrades: NVIDIA® Quadro® 5500 & GeForce™ 7950!

Xi Certified for AutoCAD® 2007, 3D Studio® Max/Viz, Softimage®, Solidworks® & Pro/E®. We custom configure each Xi system from 100+ options to match your application and your budget. Business Lease from \$ 59/month.

**@XiComputer.com** Intl: (949) 498-0858  
 Fax: (949) 498-0257  
 www.xicomputer.com  
**1-800-432-0486**



REUVEN M. LERNER

If you have been developing Web applications for a while, you might be wondering what the big deal is with Ajax.

# Ajax Application Design

Asynchronous is the operative word with Ajax, and here's what it's all about.

**During the past few months**, I've used this column to explore a number of technologies and techniques related to Ajax, the asynchronous JavaScript and XML paradigm that is the hottest thing in modern Web development. Everyone is scrambling to include Ajax on his or her sites, and for good reason. For users, Ajax applications appear more responsive and desk-top-like. For developers, Ajax is attractive because it breaks the one-page-per-click rule that has existed since the beginning of the Web, making new types of applications possible.

In an Ajax application, a click might force a complete page reload, as in a traditional Web application. But, it might instead fire an HTTP request in the background. The response to this HTTP request is handled (also in the background) by a JavaScript function, which can use the content to modify some or all of the page.

If you have been developing Web applications for a while, you might be wondering what the big deal is with Ajax. After all, it's neither new nor difficult for a JavaScript function to modify the current page via the DOM, is it? Perhaps not, but sometimes the most powerful ideas result not from fancy technologies, but from the clever combination of simple ones. HTML, HTTP and URLs were all fairly simple inventions, and they might not have gone very far on their own. But by combining them in just the right way, Tim Berners-Lee launched a revolution that continues to this day.

Just as the Web has changed the way that we view publishing and communication, Ajax has changed the way that we expect Web-based applications to work. Fortunately, working with Ajax requires only a few skills above and beyond what Web developers needed to know until now—particularly JavaScript, the DOM and CSS.

Last month, we built a small application that demonstrated the improved usability that Ajax brings to the table. As a visitor filled out the HTML form with a requested user name, a JavaScript function requested (via HTTP) a list of current user names from the server. The HTTP response contained a list of current users. By checking to see whether the newly requested user name was on that list, it was possible to tell the user in advance to choose something else.

This approach had many problems, but the two biggest ones were scalability and security. If our site becomes especially popular, we will have many registered users, so sending a complete list of user names will consume increasing amounts of CPU and bandwidth.

In addition, it is a large security risk to send all of the user names on a site to anyone who requests it. The odds are good that at least one of those users has chosen a poor password, which would make it easy to assume that person's identity. The implications of this security breach depend on your users, your application and your country. Some countries' legal systems might even see this as a prosecutable violation of database privacy laws.

So, for technical and security reasons alike, we need to find a better solution. An obvious candidate, and one we examine this month, involves sending the proposed user name to the

server via an Ajax request. The server's response will thus be a short "yes" or "no", indicating whether the browser should allow or prevent registration.

## Ajax Requests

An Ajax application consists of several parts:

1. A JavaScript function, defined in the Web page, that is invoked when a particular event happens. These event handler functions are common in the JavaScript world, even without Ajax. Before CSS, for example, it was common to use JavaScript to change the src attribute for an img tag whenever the mouse would hover over it (the onmouseover event) or move off of it (the onmouseout event). In the case of Ajax, the event handler function doesn't manipulate the DOM, but rather it sends an asynchronous HTTP request using the XMLHttpRequest object.
2. In our example application, the JavaScript function will create an XMLHttpRequest object and use it to invoke a program residing on the server. As a parameter to the request, we will send the contents of the username text field.
3. A server-side program that expects to receive the HTTP request, along with one or more parameters, and produces an appropriate HTTP response. The response theoretically may be in any legitimate MIME format, although XML, plain text and JSON (JavaScript Object Notation) appear to be the most popular choices. The server-side program will almost certainly not be written in JavaScript. You can choose the language in which you write this program, as well as the method in which it is invoked. The key is that it has access to the resources you need, such as a database, and that it can produce the output in the format you want. In this month's example application, the server-side program takes the username parameter and looks in the database to see if it is already in use. The XML that it returns will indicate its findings.
4. A second JavaScript function, also defined in the user's Web browser, that is invoked when the HTTP response is received. This callback function, as it is sometimes known, receives the HTTP response and then acts on it. Our callback routine will thus need to parse the Ajax HTTP response and then use the DOM to modify the current page as necessary.

## Improving Our Programs

Given the above list, how can we move from the simple program we wrote last month to one that will fulfill our scalability and security requirements?

When we created our simple Ajax user name-checking program in last month's column, we used two of these three elements. We created an HTML form (shown in Listing 1) that would let people register with our Web site by entering a user name, password and e-mail address.



# 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](http://www.intel.com/software/products).

**Find a reseller at** [www.intel.com/software/products/reseller](http://www.intel.com/software/products/reseller).



Listing 1.

## ajax-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">
function getXMLHttpRequest () {
  try { return new XMLHttpRequest("Msxml2.XMLHTTP"); } catch(e) {};
  try { return new XMLHttpRequest("Microsoft.XMLHTTP"); } catch(e)
{}
  try { return new XMLHttpRequest(); } catch(e) {};
  return null;
}

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);
}

var xhr = getXMLHttpRequest();

function parseUsernames() {

  // Set up empty array of usernames
  var usernames = [ ];

  // Wait for the HTTP response
  if (xhr.readyState == 4) {
  if (xhr.status == 200) {
    usernames = xhr.responseText.split("\n");
  }
  else
  {
    alert("problem: xhr.status = " + xhr.status);
  }
}

}

// Get the username that the person wants
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;
}

}

function checkUsername() {
  // Send the HTTP request
  xhr.open("GET", "usernames.txt", true);
  xhr.onreadystatechange = parseUsernames;
  xhr.send(null);
}

</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>

```

We then indicated that whenever the username text field was changed, the checkUsername JavaScript function should be invoked:

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

checkUsername then asked our server—the same server from which the current page of HTML came—for the contents of a text file:

```
function checkUsername() {  
    // Send the HTTP request  
    xhr.open("GET", "usernames.txt", true);  
    xhr.onreadystatechange = parseUsernames;  
    xhr.send(null);  
}
```

This is the first place where we will need to make a change. Rather than send a GET request without any parameters to request a static document, we will send a POST request with a single parameter (username), which will result in the execution of a server-side program.

Finally, our callback routine (parseUsernames) iterated over the list of user names that the server had sent, using the DOM to warn the user if it found a match. This is the other place where we will need to make a change. But in this case, the change will be a simplification. No longer will we need to parse through the user names sent by the server. Instead, we will need to identify only whether the response was positive or negative.

### Sending a POST Request

Last month's version of the program sent a GET request. It is possible, and even common, to send one or more parameters with a GET request. Those parameters are then stuck onto the URL, as follows: <http://www.example.com/foo.pl?param1=value1&param2=value2>.

A separate type of request, known as POST, puts the parameters inside of the request body. This has several advantages, including cleaner URLs and no limit on the length of the parameter names and values. (Many browsers limit the total size of a URL, which includes the parameters for a GET request.)

Although it is not strictly necessary for us to use a POST request for this example program, it is good to see how we can pass parameters in our request. And indeed, it is quite easy to do so. Compare the following code (taken from Listing 2) with the similar excerpt above (from Listing 1):

```
function checkUsername() {  
    // Send the HTTP request  
    xhr.open("POST", "/cgi-bin/check-name-exists.pl", true);  
    xhr.onreadystatechange = parseResponse;  
  
    var username = document.forms[0].username.value;  
    xhr.send("username=" + escape(username));  
}
```

As you can see, we have changed the first two parameters to xhr.open to be POST (instead of GET) and to point to a program that will generate dynamic output. The third parameter, which tells the XMLHttpRequest object that it should make the query in the background (that is, asynchronously), remains set to true. I also changed the name of the callback routine to parseResponse, from parseUsername.

The other change is that we are now sending parameters to the server. The variable queryString is just a string consisting of name-value pairs, in the traditional Web format of:

```
param1=value1&param2=value2
```

We thus build such a query string, and send it to the server.

**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™- (Rev.F) Starting @ \$5293.20

1 of 4 AMD Opteron 8212 CPU  
12 GB DDR2 ECC REG PC-3200  
1 of 3 73GB SCSI Drive  
2 GigaE, CD, FD,  
Optional Remote Management Card (IPMI)  
Call for more choices of AMD socket F Servers.



#### 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



#### 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.

Listing 2.

**post-ajax-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">
function getXMLHttpRequest () {
  try { return new XMLHttpRequest(); } catch(e) {};
  try { return new XMLHttpRequest(); } catch(e)
{}
  try { return new XMLHttpRequest(); } catch(e) {};
  return null;
}

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);
}

var xhr = getXMLHttpRequest();

function parseResponse() {
  // Get variables ready
  var response = "";
  var new_username = document.forms[0].username.value;
  var warning = document.getElementById("warning");
  var submit_button = document.getElementById("submit-button");

  // Wait for the HTTP response
  if (xhr.readyState == 4) {
  if (xhr.status == 200) {
    response = xhr.responseText;

    switch (response)
    {
      case "yes":
        setText(warning,
          "Warning: username '" +
            new_username + "' was taken!");
        submit_button.disabled = true;
        break;

      case "no":
        removeText(warning);
        submit_button.disabled = false;
        break;

      case "":
        break;

      default:
        alert("Unexpected response '" + response + "'");
    }
  }
  else
  {
    alert("problem: xhr.status = " + xhr.status);
  }
}

function checkUsername() {
  // Send the HTTP request
  xhr.open("POST", "/cgi-bin/check-name-exists.pl", true);
  xhr.onreadystatechange = parseResponse;

  var username = document.forms[0].username.value;
  xhr.send("username=" + escape(username));
}
  </script>
</head>
<body>
  <h2>Register</h2>
  <p id="warning"></p>
  <form action="/cgi-bin/register.pl" method="post"
  enctype="application/x-www-form-urlencoded">
    <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>

```

## The Server Side

Ajax is almost exclusively a client-side paradigm. And, indeed, it is increasingly clear that we can use JavaScript in general, and Ajax in particular, to create new and interesting applications and interfaces. That said, server-side programs still have a major role to play in Web applications, including Ajax applications.

To begin with, only server-side programs can access the site's relational database. (And yes, it's theoretically possible to have JavaScript access the database directly, but that would be a security and performance nightmare.) This means everything you normally would store in a database, but want to have displayed in the browser, will need to be filtered through a server-side program. Almost any nontrivial application will thus benefit from being part of a larger Web framework, such as Zope, Ruby on Rails or even a roll-your-own system that encapsulates behavior in a set of related methods or functions. In other words, the server-side programs in an Ajax application become very specialized database query and reporting tools.

In the interests of time and space, we don't access a database this month. However, there is no way for the HTTP client to know whether the HTTP server is checking a database or returning a random result, and we will take advantage of this secrecy to fudge the lack of a database. If we decide at some point to modify our server-side program to retrieve a list of user names from a database instead of hard-coding the list in a hash, that will be just fine.

Our server-side program, `check-name-exists.pl` (Listing 3), is a simple CGI program written in Perl. We turn the `POSTDATA` parameter, which we have received from the Ajax request, and look inside it to see if we received a setting for username. If so, we then look for a match among the keys of the `%usernames` hash. If we find a match, it returns yes to the caller. If there is no match, it returns no.

Notice how we use a hash, rather than an array, to store the user names. This is a hack for the sake of efficiency; the time it takes to find an array element (and see if there is a match) is proportional to the number of elements in the array. By contrast, hash key lookups take constant time, regardless of how many elements there are. In a production setting, we obviously would expect to look for user names in a database or server-side disk file, rather than a hash or an array.

This example also demonstrates one way to mock up an Ajax application while development is still taking place—create a server-side program that produces results for a very small subset of the data, simulating the full range of database queries that you might normally want to use. In this way, development

[www.faircom.com/go/?track](http://www.faircom.com/go/?track)

Your packages get

# tracked & delivered

on time by FedEx.

FairCom database technology makes it possible.



**FairCom**®

on the JavaScript side of the project will not have to wait for the server-side portion to be complete, allowing for more parallelized development.

### Parsing the Response

When the response arrives from the server, our callback routine, `parseResponse`, is invoked. As always, we wait until the `readyState` of our `XMLHttpRequest` is 4 and for the HTTP status code to be 200. At that point, we can expect one of four different responses from the server:

- A yes response indicates that the user name was taken. We disable the form's submit button and display a warning. If and when the user changes the text inside

of the username text field, the warning will be removed and the submit button re-enabled.

- A no response indicates that the user name is available. We remove any warning that might have been placed, and enable the submit button.
- An empty response might come before the yes or no, in which case we ignore it.
- Finally, it's possible that our program will not behave precisely as we might expect. If this happens, we display the unexpected response that we received for debugging purposes. This is the sort of thing you would probably want to remove from production code.

Listing 3.

#### check-name-exists.pl

```
#!/usr/local/bin/perl

use strict;
use diagnostics;
use warnings;

use CGI;
use CGI::Carp;

# Define the usernames that are taken
# (Use a hash for lookup efficiency)
my %usernames = ('abc' => 1,
                'def' => 1,
                'ghi' => 1,
                'jkl' => 1);

# -----
my $query = new CGI;
print $query->header("text/plain");

# Get the POST data
my $postdata = $query->param("POSTDATA");

# Get the username
my ($name, $value) = split /=/, $postdata;

my $username = '';
if ($name eq 'username')
{
    $username = $value;
}

# If this username is defined, say "yes"!
if (exists $usernames{$username})
{
    print "yes";
}

# Otherwise, say "no"!
else
{
    print "no";
}
```

Notice how we used a switch statement to look at the different possibilities. Also notice how we were able to reduce the complexity of our JavaScript code by sharing the work with the server. This is the key to a good Ajax application. Rather than having the client or the server do all of the work itself, each of them shares in the burden, doing what it can do fastest and most cleanly.

Finally, you might notice that for all of our talk about XML—it is, after all, the x in Ajax—there was a distinct lack of XML in this application. True, we used the `XMLHttpRequest` to send HTTP requests to the server, but what happened to the XML?

The truth is that Ajax is a great name, but it doesn't quite describe the range of options the programming paradigm provides. The HTTP response, as I indicated above, can come in any MIME type, although XML and plain text are the most common. If this application were returning a more sophisticated set of data, such as a store inventory or points for a chart, XML might be more appropriate. Another format that is gaining in popularity is JSON, which resembles Perl's "Data::Dumper" in its representation of JavaScript objects. Ajax is merely a technique for dividing the work between the client and the server; you should not feel compelled to use XML for the data transfer if it is inappropriate for the task at hand.

### Conclusion

This month, we finally produced an application worthy of the Ajax moniker. We used a combination of JavaScript (on the client side) and Perl (on the server side) to check whether a user name was already taken. In doing so, we saw how to use the POST method for submitting data and sent a named parameter to the server. In making these changes, we turned a simple, insecure and unscalable program into a relatively secure and scalable one, without sacrificing the immediate response and interactivity that Ajax brings to the table.

At the same time, you might have noticed our HTML page contained a large number of functions that will be useful for a wide variety of Ajax applications. Starting next month, we will look at some of the open-source libraries that make it easier to create Ajax applications, allowing you to concentrate on the higher-level details. ■

---

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



# 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](http://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.



MARCEL GAGNÉ

# Tonight's Menu: Diner's Choice!

The December 2006 Cooking with Linux marks seven years of *Chez Marcel*, its Linux Chef and François, his famous waiter. For this occasion, Marcel has invited some regular readers of the column to share their favorite issues and, of course, a glass or two (or three, or four) of wine.

**Seven years, mon ami.** Yes, François, it does seem to have gone by very quickly. We'll have plenty of time to reminisce when our guests arrive. Ah, but they are already here! Welcome, *mes amis* to *Chez Marcel*, where you'll find one of the world's largest wine cellars and the finest in Linux and open-source software. Please sit and make yourselves comfortable. I'll have François fetch the wine right away.

*Quoi? Incroyable!* It appears that my faithful waiter decided to beat me to the punch when it came to tonight's wine selection. He has already chosen the wine and brought it up from the cellar. Since this is a penguin-studded magazine, he has chosen one of my favorite, not to mention inexpensive, Australian wines. It's called Little Penguin. Of course, my judgment may be clouded by the name of the winery, but the choice seems fitting. François, please serve the Little Penguin wine for our guests. Today, we have both a Chardonnay and a Shiraz (for those who prefer red).

Those of you who have been following this column from

the beginning may already be aware that this December 2006 issue marks seven years of my Cooking with Linux. The first of these columns, however, was a kind of experiment, featured in the September 1999 special issue. The regular series began a few months later, with the January 2000 issue. To commemorate, the folks at *Linux Journal* suggested I select a few of my favorite columns for this issue. When I thought about this idea, I decided that the best arbiters were the readers, and so I've invited a number of special guests here tonight to tell us what they liked, and why.

Let's start with table seven, where Troy Banther is waving his hand madly. "One of my favorites is 'If Only You Could Restore Wine' in the June 2006 issue of the magazine. Knowing what programs are out there in the Linux world and having a how-to on backing up and restoring is great. I'm extremely partial to KDar since I use the KDE window manager."

Jon Biddell, at table 15, agrees with Troy Banther. "I think my favorite would have been the June 2006 article on backups—something I don't do anywhere near as often as I need to...."

Over at table three, Colleen Beamer says she doesn't know if she can pick just one. She tends to lean toward the August 2004 issue because "this column introduced me to Krecipes [Figure 2], and I went from there to have

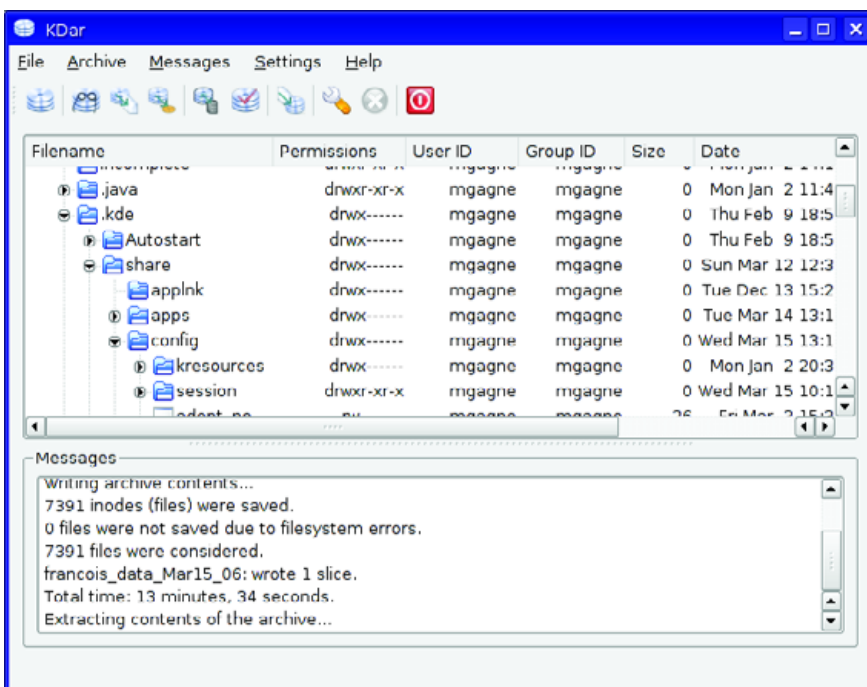


Figure 1. KDar, because backups should be easy, even if they are essential.

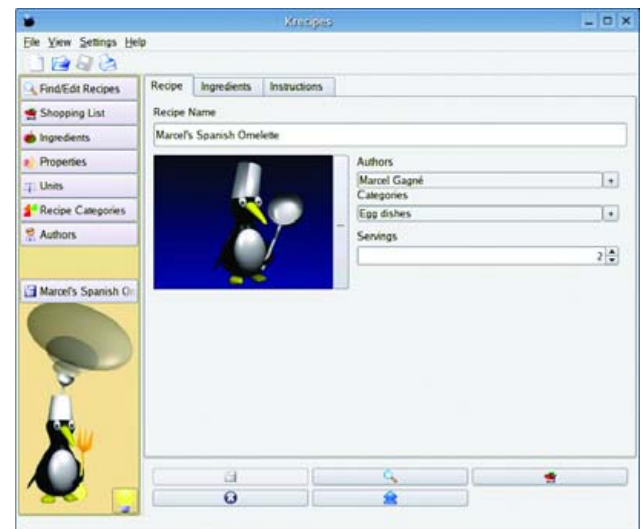



Figure 2. Krecipes as it appeared in August 2004. The package will likely reach its 1.0 release when this issue is released.

Provides the flexibility and scalability for demanding enterprise-class applications. The AMD Opteron™ processor's forward-thinking technology provides flexibility with a 1-8P scalable design (up to 16 cores), delivering the headroom needed for current high-end computing needs or future applications and technologies.



AMD Opteron™ Processor with Direct Connect Architecture Performance

SPECfp®\_rate2000 Linux OS



**ZT Reliant Server A8200**

- Dual-core AMD Opteron™ Processor Model 2210 (1.80GHz, 2x1MB L2 Cache)**  
 (Upgradeable to 2 X Dual-Core AMD Opteron™ 2000 Series Processors )
- 1GB ECC Registered DDR2 533MHz Memory (Supports up to 32GB)
  - 2x 250GB SATA2 16MB Cache Enterprise Hard Drive (RAID 1)
  - 8x 1" Hot-Swappable SATA/SAS Drive Bays
  - ATI® ES1000 16MB Graphics Controller
  - nVidia® MCP55 Pro SATA2 3.0Gbps Controller (RAID 0, 1, 0+1, 5, JBOD Support)
  - DVD-ROM and Floppy Drive
  - nVidia® MCP55 Pro Chipset Dual-port LAN / Ethernet Controller
  - 4U Rackmount/Tower Chassis w/ 650W Redundant-Cooling Power Supply
  - 8x Fan Support with Speed Control
  - Super Doctor III Server Management Software
  - 3-Year Limited Warranty

**\$1,899**

**ZT Reliant Server A8201**

- Dual-core AMD Opteron™ Processor Model 2212 (2GHz, 2x1MB L2 Cache)**  
 (Upgradeable to 2 X Dual-Core AMD Opteron™ 2000 Series Processors )
- 1GB ECC Registered DDR2 533MHz Memory (Supports up to 32GB)
  - 2x 250GB SATA2 16MB Cache Enterprise Hard Drive (RAID 1)
  - 8x 1" Hot-Swappable SATA/SAS Drive Bays
  - ATI® ES1000 16MB Graphics Controller
  - nVidia® MCP55 Pro SATA2 3.0Gbps Controller (RAID 0, 1, 0+1, 5, JBOD Support)
  - Slim DVD-ROM and Floppy Drive
  - nVidia® MCP55 Pro Chipset Dual-port LAN / Ethernet Controller
  - 2U Rackmount Chassis w/ 700W Redundant Power Supply
  - Super Doctor III Server Management Software
  - 3-Year Limited Warranty

**\$2,399**

**ZT Reliant Server A8202**

- Dual-core AMD Opteron™ Processor Model 2216 (2.40GHz, 2x1MB L2 Cache)**  
 (Upgradeable to 2 X Dual-Core AMD Opteron™ 2000 Series Processors )
- 2GB ECC Registered DDR2 533MHz Memory (Supports up to 32GB)
  - 4x 500GB SATA2 16MB Cache Enterprise Hard Drive
  - 12x 1" Hot-Swappable SATA/SAS Drive Bays
  - ATI® ES1000 16MB Graphics Controller
  - nVidia® MCP55 Pro SATA2 3.0Gbps Controller (RAID 0, 1, 0+1, 5, JBOD Support)
  - DVD-ROM and Floppy Drive
  - nVidia® MCP55 Pro Chipset Dual-port LAN / Ethernet Controller
  - 3U Rackmount Chassis w/ 800W Redundant Power Supply
  - Super Doctor III Server Management Software
  - 3-Year Limited Warranty

**\$3,999**

Get More Than You Expected!

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

Go to [ztgroup.com/go/linuxjournal](http://ztgroup.com/go/linuxjournal)

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

Promote code: ij1206



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

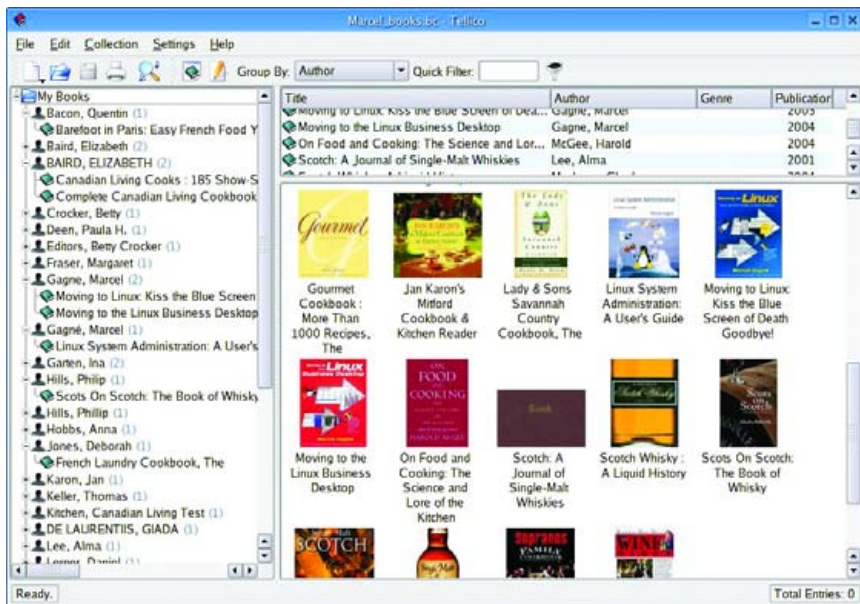


Figure 3. Tellico makes a great personal library system, and it looks good doing it.



Figure 5. The Ultimate in Take-Anywhere Linux

Tellico [Figure 3], in the April 2005 column titled, 'The Cook's Collection'. Why? Simply because people are always borrowing my books."

He adds, "*En passant, un article sur le caractère multilingue de Linux pourrait être intéressant, n'est-ce pas?*"

All right, Daniel, I'll just take care of that request right now. In KDE, fire up the KDE Control Center (command name, `kcontrol`) and click on Regional & Accessibility. Under that category, select Keyboard Layouts. Usually, the only layout visible to the right in the main window (under Active Layouts) is whatever you chose during installation. On mine, it says, "U.S. English". Now, I like having access to a quick keyboard switch so that I can use things like the `é` at the end of my last name. These are included in the Canadian layout. I also can get them using the "U.S. English w/deadkeys" layout. So, I select it from the Available Layouts section, then click the Add button (Figure 4).

If you want more, add them now as well. Then, click Apply. A small flag icon appears in the system tray of your kicker panel, over on the right. Click the icon to switch from one layout to the other. If you are using OpenOffice.org, you now can type the characters you want without doing an "Insert Special Character" operation each time. If you find yourself having problems with those characters displaying properly in OpenOffice.org, this is usually a problem due to using incorrect or incompatible fonts. For example, someone sends you a document with a Microsoft font and you don't have it installed. The resulting text, particularly if you are entering special characters, can then look a bit strange. All that's needed is to add those fonts using the `spadmin` program (OpenOffice.org's printer administration program, which also lets you add fonts). I hope that helps you out, Daniel.

François, our guests' glasses are looking a little low. If you would be so kind as to visit everyone and indulge each guest with his or her choice of wine. While you do so, we'll check in with John Kerr, over at table 12. He says, "I must admit that finding a favorite column would be difficult as there are so many good ones to choose from. My favorite column, howev-

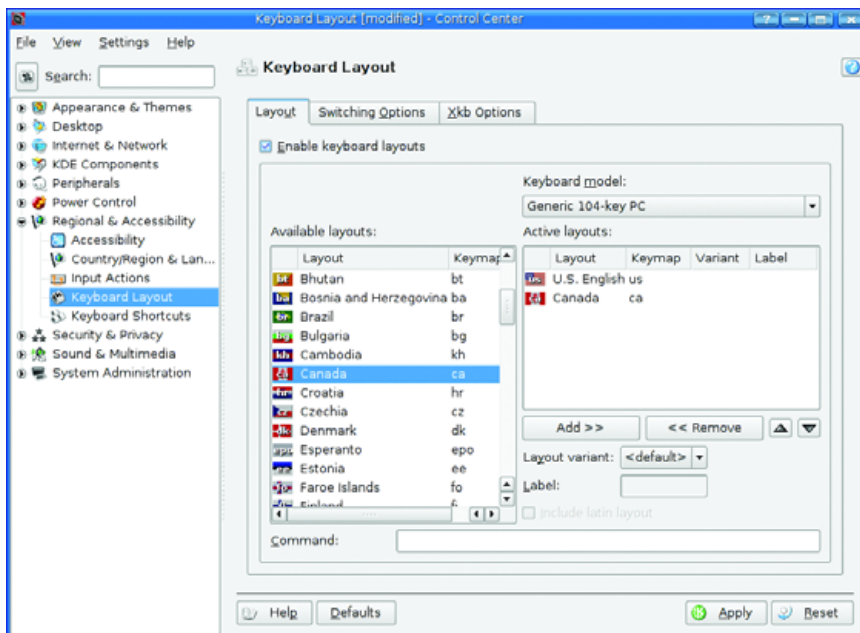


Figure 4. You don't have to be limited to a single keyboard layout.

my first true open-source involvement by writing the *Krecipes Handbook*. 'Crossing Platforms', in the May 2005 issue, introduced me to my *still* favourite Linux game, *Blobwars*. However, if I have to pick one column, the one in the September 2005 issue, 'Wireless Tools', has helped me the most. Without it, I probably wouldn't have been able to know what to do to get wireless installed and working on my laptop."

When Daniel Gagnon, who is sipping his wine at table 27, was asked about his favorite, he replied, "The one about



Figure 6. IceWM Nested inside of GNOME Nested inside of KDE



Figure 7. Video conferencing with VIC—who is this mysterious woman, chatting with our François?

er, is the August 2005, 'Ultimate in Small Linux' article. This column demonstrates two mini-distros (Damn Small Linux and Puppy Linux) as well as the unique ways they can be used—live CD, disk install or USB key [Figure 5]. These distros also can be used to revive older equipment. This only emphasizes the diversity of our favorite operating system."

Choosing a favorite was also a bit difficult for David Knickmeyer, sitting over at table 19. He says, "I'd probably go with the Xnest article [Figure 6], 'Can't Get Enough Desktops', in the March 2004 issue. I admit it, I'm an eye-candy junkie. GNOME, Enlightenment, Window Maker—I've tried them all, often. I keep coming back to KDE and SuperKaramba, but the nested X servers let me play around and still be productive—well...."

I see another hand waving over at table eight, and it's Lew Pitcher. "My list runs around 20 articles so far, and I've covered only about half of my LJ back issues. So, you can imagine the difficulty I have in picking one col-

umn. But, pick one I have. I take you back to February 2002, to a column entitled 'Observe, Mon Cher Ami', in which you introduced us to xawtv and video conferencing in general. In that column, François finally got a face (even though it was a stuffed toy avatar), and he met a mystery woman, who shall go nameless, but might be known as Sally in other places [Figure 7]. Of course, my second-favorite column would be from December 2004, 'Lights...Camera...Action!', in which we find the tools to do video podcasts among the various bits and pieces of our toolboxes. Who knew that a cheap microphone and a just-as-cheap Webcam could make you a video star?"

Over at table 32, Johann Schmidt has this to say, "My nomination for favorite Cooking with Linux column is fairly recent—from the December 2005 issue, 'amaroKing the Night Away'. I had not heard of AmaroK before [Figure 8], and your introduction to it was fantastic. I installed it at home the day I received the issue at work, and after only a slight bit of dependence-finding, I had an awesome music system/library/jukebox. It works with our iPods too! I have been getting my children more and more exposed to Linux—and because AmaroK doesn't crash like iTunes does on that proprietary x86 OS, they have yet another reason to be open-source kids."

Bert Sutherland has this to say, "My favorite choice by far is your November 2004 'Illuminating Your Network's Darkest Corners'. I had been a *Linux Journal* subscriber for only a few months when this issue came out, and it is still one of the best articles I've read. When I sit down to read *Cooking with Linux*, I know I am in for a surprise and that you will be

## LINUX ONLINE TRAINING

Real Training on Real Servers

**Basic Linux Administration Course - \$295.00**

**Network Intrusion Course - \$295.00**

**Ubuntu Server Course - \$295.00**

**SpiderTools.com**

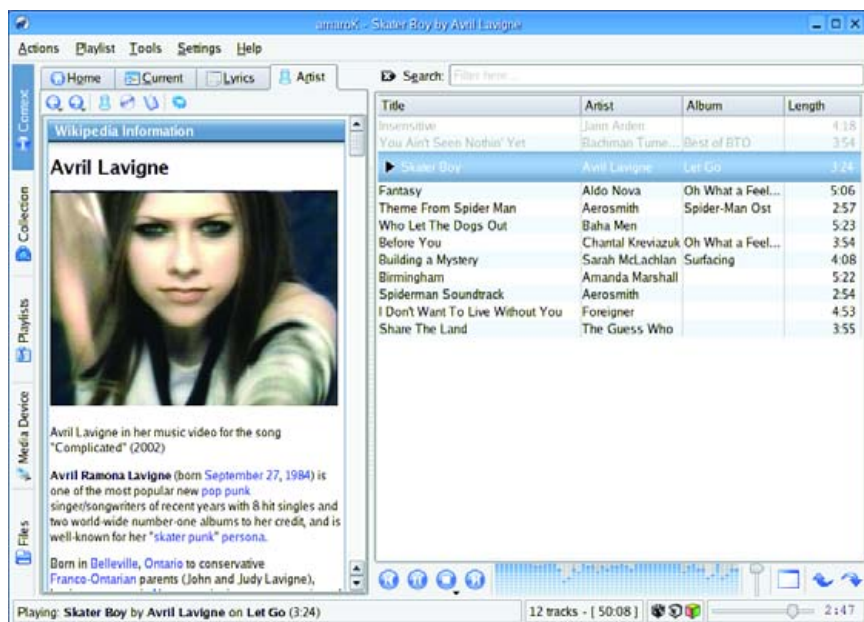


Figure 8. Amarok, one of the best music players out there, just keeps on getting better all the time.




Figure 9. IPTraf's Default Monitoring Window

François was designed by the amazing Robert Karlsson, courtesy of *Linux Journal*. My thanks to both for putting a face (and species) on my faithful waiter.

servicing up another selection (of wine *and* Linux tidbits of information) that I was previously not aware existed. This was the case with the programs mentioned in your network article, and to this day IPTraf [Figure 9] is one of the first programs I install when working with a new machine."

Margaret Wendall, sitting over at table 14, also is interested in security: "It's become apparent to me and one of my clients (I help with his Web pages) that



Figure 10. Using the Steghide program, a rather large order of wine is hidden and encrypted within this portrait of François.

people are stealing *my* photos and using them to advertise on their pages. I think your good instructions on steganography, back in the January 2005 article 'Forgotten Security', are going to be a better solution than something like Digimark, because only two or three people will even know the photos are encrypted. This is not only fun, but very practical." *Was I a spy in a past life?*

François, how did you manage to sneak in that portrait of yourself? You truly are efficient today, *mon ami*. I see that you have continued to refresh our guests' glasses. Excellent. Sadly, it seems that closing time is upon us, so the next refill will have to be the last one. My sincere thanks to everyone out there for joining me here every month these past seven years and for helping make Cooking with Linux as much fun as it has been. I also want to thank the members of my own WFTL-LUG, aka "The Lug Nuts" (new members welcome), who joined me in the restaurant today. For a complete list of past Cooking with Linux columns and links to each article on-line, check out [www.marcelgagne.com/ljcooking.html](http://www.marcelgagne.com/ljcooking.html). Now that François has so graciously refilled your glasses, please join me in a toast and let us all drink to one another's health. *A votre santé! Bon appétit!* ■

**Resources for this article:** [www.linuxjournal.com/article/9379](http://www.linuxjournal.com/article/9379).


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 [mggagne@salmar.com](mailto:mggagne@salmar.com). You can discover lots of other things (including great Wine links) from his Web site at [www.marcelgagne.com](http://www.marcelgagne.com).

# T20DC-1U


## Custom Chassis

A Division of  
**M&A**  
Technology

**Socket F  
Opteron™  
Tested!**



**Tested with Dual Core Opterons**



**Fan Assembly for Maximum Performance**



**1U Density for Scalability**




**Easy Access to I/O and Power**

### Highlights

- Custom 1U Density Design
- Dual-Core Opteron Tested
- "Socket F" ready
- Top rated Thermals
- Fits standard 19" Rack
- Power & Reset buttons
- Power & HDD LED
- Front USB
- 4x Front 40mm Fans
- 6x Middle 40mm Fans
- 2x Rear 40mm Fans
- Supports SATA, SAS and SCSI HDD
- Includes 1x 5.25" Bay
- Support for PCI-Express, HTX and PCI-X Risers
- 500W Power Supply
- Convenient Front Handles
- Slide Rails Included

**The 1st InfiniPath  
Interconnect Cluster**



**UC-Davis - based on  
the Dual-core Opteron™**

**TeamHPC**  
A division of  
**M&A Technology, Inc.**  
**1-866-TEAMHPC**

[www.teamhpc.com](http://www.teamhpc.com)



DAVE TAYLOR

Unfortunately, the `expr` command that we're using for the mathematical calculations doesn't have the ability to work with these powers of two, so we're going to have to do the work ourselves, converting massive numbers into more readable KB, MB or GB values, as appropriate.

# Breaking Numbers Down

A kilo of information on how to represent even giga numbers in a mega-useful way.

**Last month**, we continued our journey into the dark caverns of Apache Web logs, examining how relatively simple shell scripts can be utilized to produce useful and important data. The specific script we created searched a log file for traffic that occurred the previous day, summarizing the number of bytes transmitted.

That's all well and good, but as with many shell scripts, there's a bit of a problem with this one, which was immediately obvious when my busy site produced an estimated monthly data transfer rate of 2346990660 bytes.

Clearly that's a very human-unfriendly number, and doubly so without any commas to break it up into thousands, millions and so on. More important, when talking about data transfer, we're used to thinking in terms of powers of two, so 1 kilobyte is 1024 bytes of data, not 1000 bytes of data, and 1 megabyte is 1024 kilobytes of data, and so on.

Unfortunately, the `expr` command that we're using for the mathematical calculations doesn't have the ability to work with these powers of two, so we're going to have to do the work ourselves, converting massive numbers into more readable KB, MB or GB values, as appropriate.

## Converting Numeric Values

The basics are pretty easy:

```
kilo=$(( $value / 1024 ))
mega=$(( $kilo / 1024 ))
giga=$(( $mega / 1024 ))
```

Given a nice huge number like 2346990660, the results are then quickly calculated:

```
$ sh -x convert.sh 2346990660
+ value=2346990660
+ kilo=2291983
+ mega=2238
+ giga=2
+ exit 0
```

(Helpful tip: the `-x` option lets you debug shell scripts by showing, line by line, what command is being executed.)

The problem with this approach is immediately obvious when we switch from a huge number, more than 2GB, to a smaller value:

```
$ sh -x convert.sh 5000
+ value=5000
+ kilo=4
+ mega=0
+ giga=0
+ exit 0
```

We don't want zero values; we want to see the fractional decimal values, which means not only that we can't use the built-in mathematical capabilities of the shell, but we also can't

use `expr`. Instead, we need to move into the crafty, ancient world of `bc`, the binary calculator.

Now, `bc` isn't for the faint of heart, but to save you from reading the man page, here's how you can force four digits after the decimal point on the result of a division that results in a value less than 1.0:

```
$ echo "scale=2 ; 3000 / 30001" | bc
.0999
```

Can you see how to put these together? Here's a new, far-improved way to calculate kilo, mega and giga:

```
$ sh -x convert.sh 5000
+ value=5000
++ echo 'scale=2; 5000 / 1024'
++ bc
+ kilo=4.88
++ echo 'scale=2; 4.88 / 1024'
++ bc
+ mega=0
++ echo 'scale=2; .00 / 1024'
++ bc
+ giga=0
+ exit 0
```

The debug output from the `-x` option is getting a bit confusing here, I admit, but you now can see that kilo is set to 4.88 when given the initial value of 5000 bytes, and that both mega and giga are zero.

Let's try again (and I'll clean up some of the spurious debug output from this point on, for clarity) with the initial really big value:

```
$ convert.sh 2346990660
value=2346990660
kilo=2291983.06
mega=2238.26
giga=2.18
```

Cool. Now we can finally see that we're talking about 2.18GB of data being transferred off the site each month—far more coherent than the huge value shown earlier.

Now, let's figure out how to show always the most logical of these values, rather than all of them.

## Displaying the Simplest Answer Only

The easiest way to figure out which value is best is simply to ascertain where the value drops below 1.0. In the case of 5000 bytes, that'd be best displayed as 4.88KB, and in the case of the big number, that's 2.18GB.

To figure out when the value drops below zero, we'd love to have a floating-point numeric comparison, but sadly, the shell can't manage it. If you try it, you'll just get the error "integer expression expected".



### FOCUS ON WHAT MATTERS MOST

GET DUAL-CORE INTEL XEON PROCESSORS IN SERVERSDIRECT SYSTEMS AND CONCENTRATE ON BUSINESS OPPORTUNITIES INSTEAD OF SERVER HEADACHES



**STARTING PRICE**  
**\$2,499**

## SDR-3500T 3U 64-BIT XEON DATABASE SERVER

Designed to provide outstanding performance with dual-core Xeon® 5000/5100, ideal solution for storage and mission critical business application server

### FEATURE HIGHLIGHTS:

Intel Xeon 5050 3.0GHz Dual Core Processor (Dual Processor Option)

3U Chassis with 800W Redundant Power Supply

Intel® 5000P (Blackford) Chipset

Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2pcs x 512MB)

Seagate 400GB SATA-II 16MB Cache 7200RPM

16 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

### 1U ENTRY LEVEL SERVER

1U Pentium D servers support the latest dual-core architecture that delivers unmatched multi-tasking capabilities in an uni-processor environment

1U Rackmount Chassis with 300W Power Supply

Intel® Pentium® D 930 3.0GHz Dual Core 800FSB

Intel® E7230 (Mukilteo) Chipset

Kingston 1024MB 667MHz DDR2 ECC (2pcs x 512MB)

Seagate 250GB SATA-II 7200RPM hard drive

2 x 1" Hot-swap SATA Drive Bays

ATI RageXL graphics, 8MB

2x Intel® 82573V (Tekoa) PCI-e Gigabit LAN Ports

RAID 0, 1

**SKU# SDR-1110T**



**\$999**

### COST EFFECTIVE 2U SERVER

Businesses with 50 or more employees seeking maximum performance, storage room, expandability and uptime

Intel Xeon 5050 3.0GHz Dual Core Processor (Dual Processor Option)

2U Chassis with 600W Power Supply

Intel® 5000V (Blackford VS) Chipset

Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2pcs x 512MB)

Seagate 400GB SATA-II 16MB Cache 7200RPM

6 x 1" Hot-swap SATA-II Drive Bays

ATI ES1000 Graphics with 16MB video memory

Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller

RAID 0, 1, 5, 10 support

**SKU# SDR-2501T**



**\$1,599**

### 4U TOWER/RACKMOUNT 64-BIT XEON ENTERPRISE SERVER

Powered by the latest dual-core Xeon 5000/5100® sequence processors, versatility in either a tower or rackmount 4U form factor and support enterprise-class

Intel Xeon 5050 3.0GHz Dual Core Processor (Dual Processor Option)

4U Rackmountable / Tower

Intel® 5000X (Greencreek) Chipset

Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2pcs x 512MB)

8pcs x Seagate 400GB SATA-II 16MB Cache 7200RPM

3Ware 9550SX-8port RAID Controller Card

GeForce 6200 TC 128MB OB 256M support DDRTV-out PCI-E

Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller

**SKU# SDP-7045A-TB**



**\$3,999**

### SDR-5500T 5U ADVANCED STORAGE SERVER

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

Intel Xeon 5050 3.0GHz Dual Core Processor (Dual Processor Option)

5U Chassis with 24 hot-swap bays & 950W redundant power supply

Intel® 5000P (Blackford) Chipset

Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2pcs x 512MB)

12pcs x Western Digital 400GB SATA RAID Drive

1pc x 3Ware 9550SX-12 port RAID Controller Card

ATI ES1000 Graphics with 16MB video memory

Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller

RAID 0, 1, 5, 10 support

**SKU# SDR-5500T**



**\$6,799**

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

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

**1.877.727.7887 | www.serversdirect.com**



There are a number of ways to get the “floor” of the value, but I use bc again here to do the job by calculating the division once more, this time without any scale value at all:

```
kiloint=$( echo "$value/1024" | bc )
```

Doing this gets just the integer portion of the \$kilo value, and that can indeed be tested in a conditional statement:

```
if [ $kiloint -lt 1 ] ; then
```

Now, put it all together, and here’s how the script looks:

```
kilo=$( echo "scale=2; $value / 1024" | bc )
kiloint=$( echo "$value / 1024" | bc )
```

```
mega=$( echo "scale=2; $kilo / 1024" | bc )
megaint=$( echo "$kilo / 1024" | bc )
```

```
giga=$( echo "scale=2; $mega / 1024" | bc )
gigaint=$( echo "$mega / 1024" | bc )
```

```
if [ $kiloint -lt 1 ] ; then
  echo "$value bytes"
elif [ $megaint -lt 1 ] ; then
```

## RUGGED EMBEDDED SERVER



- Intel ULV Celeron 400MHz CPU
- 256MB RAM Expandable to 512MB
- 256MB CompactFlash™
- Dual 10/100 Base-T Ethernet
- Reliable (No CPU Fan or Disk Drive)
- Three RS-232 & One RS232/422/485 Serial Ports
- 6 General Purpose I/O Lines
- DC input integrated power supply



2.6 KERNEL

**Rugged SIB**  
(Server-In-a-Box)  
Starting at \$750.00  
Quantity 1.

- EMAC Linux 2.6 Kernel
- Menu Drive Configuration Utility
- Eclipse Development Environment
- HTTP and FTP Servers
- PPP Dial In/Out Server & Client
- Telnet Server

Since 1985  
OVER  
21  
YEARS OF  
SINGLE BOARD  
SOLUTIONS

**EMAC, inc.**  
EQUIPMENT MONITOR AND CONTROL

Phone: (618) 529-4525 • Fax: (618) 457-0110 • [www.emacinc.com](http://www.emacinc.com)

```
  echo "${kilo}KB"
elif [ $gigaint -lt 1 ] ; then
  echo "${mega}MB"
else
  echo "${giga}GB"
fi
```

A little funky, but it certainly works exactly as we’d hope:

```
$ sh convert.sh 5000000000
4.65GB
$ sh convert.sh 5000000
4.76MB
$ sh convert.sh 50000
48.82KB
$ sh convert.sh 50
50 bytes
```

The final step is to make it a function so we can include it in other shell scripts and access it as desired. This is done within the Bourne Shell by giving it a unique name and then wrapping the functional code in braces:

```
kmg()
{
  code for function goes here, params are $1, $2, etc.
}
```

This can then be invoked within a shell script by name (k=kilo, m=mega, g=giga):

```
kmg 5000000
```

More important, you can embed it within a line by using a subshell notation, so given the kmg() function, the following two-line script works splendidly:

```
echo given value is $1
echo which converts to $(kmg $1)
```

That’s nice and short, and if the kmg function is dropped into its own file, you also can use the . command to include another file in the shell script, meaning that the entire test script is now:

```
#!/bin/sh
. kmg.sh

echo The given value $1 bytes = $(kmg $1)

exit 0
```

I’m out of space here, but I hope you can see how this approach can be applied to a wide variety of different shell tasks, making your shell scripts far more efficient and faster to write too! ■

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](http://www.intuitive.com).

# SUPERMICR<sup>®</sup>

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

### SuperServer 6015X-8/T

- Dual Intel<sup>®</sup> Dual-Core 64-bit Xeon<sup>™</sup> 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/I<sup>2</sup>C 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<sup>®</sup>** 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.arrownaep.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



MICK BAUER

As far as I'm concerned, when using UML in security scenarios, COW files are mandatory.

# Running Network Services under User-Mode Linux, Part II

Populate and network your very own virtual network server.

Here in the **Paranoid Penguin column**, we're in the midst of building a virtual network server using User-Mode Linux. Last month, I explained why this is a good idea, how it works, how to prepare your host for optimized User-Mode Linux operation and how to build a kernel for your guest (virtual) system(s).

This month, we turn our attention to the guest system: how to obtain a prebuilt root filesystem image, how to configure networking on both your host and guest systems, and how to begin customizing the root filesystem image for your own purposes.

## Quick Review

First, here's a quick review of what we're trying to do, in case you missed last month's column. Our objective is to use User-Mode Linux to create one or more virtual guest machines, each running a different network service. That way, if one application (for example, BIND) on one guest machine gets compromised somehow, Sendmail, Apache and whatever else you've got running on other guest systems (or on the underlying host system itself) won't be affected.

(Per User-Mode Linux convention, we're using the word *host* to denote a system on top of which virtual machines run and the word *guest* to denote a virtual system instance.)

Debian is our somewhat arbitrary choice here for both host and guest systems, due to the ease with which you can create bare-bones Debian installations, though User-Mode Linux itself is decidedly distribution-agnostic. We'll create a single guest system, running BIND software for DNS services.

On the strength of last month's procedures, hopefully you've got a skas-enabled host kernel and a guest kernel compiled for the um architecture. Now, it's time to acquire or build a root filesystem image.

## Just What Is a Root Filesystem Image, and How Will It Be Used?

When your Linux host starts up, it learns where */* is via the root command-line switch; somewhere in *lilo.conf* or *menu.lst* is a kernel-invocation line containing something like `root=/dev/hda1`. That's how it works with User-Mode Linux too, except that rather than a physical hard disk, such as */dev/hda*, we generally use a virtual disk in the form of a single flat file, called a root filesystem image.

The root filesystem image contains a complete Linux distribution. You've already created similar image files yourself if you've ever copied a CD-ROM to an ISO file (or vice versa). Using a filesystem that takes the form of a single file has two important ramifications for User-Mode Linux: first, it helps keep your guest system relatively compact and portable; second, it

makes change control as simple as tracking changes to a single file, via the COW file method.

Suppose I start a User-Mode Linux guest with this command:

```
umluser@host:~> ./guestkernel ubd0=mycow,my_root_fs root=/dev/ubda
```

Note the `umluser@host` prompt. I'm executing this command from a shell session to which I'm logged in as a regular user, *not* root. `guestkernel` is my executable User-Mode Linux guest kernel; `ubd0` is a virtual disk device I'm declaring to consist of the image file `my_root_fs` plus a change-on-write (COW) file called `mycow`. The root switch defines our root partition to be the virtual disk `ubda` (identified by its full path, `/dev/ubda`).

Somewhat confusingly, by convention, virtual disk declarations use numbered device names (`ubd0`, `ubd1` and so on), but root filesystem definitions use the corresponding letters instead (`ubda`, `ubdb` and so on), which are synonymous. The command `./guestkernel ubda=mycow,my_root_fs root=/dev/ubda` actually works just as well on my SUSE system as the above command, but your distribution of choice may behave differently.

Strictly speaking, the COW file is optional. If you specify one, changes you make to the image file during your User-Mode Linux session will be written to the COW file rather than to the disk image itself. If you omit the COW filename, the image file will be written to directly by the guest kernel—that is, any changes you make to your guest system will be “permanent”.

As far as I'm concerned, when using UML in security scenarios, COW files are mandatory. A key assumption in using User-Mode Linux for hosting a network service is that this service may be compromised in some way, and if it is, you'll want to be able to recover as quickly as possible. If you use a COW file, all you'll need to do to restore a guest system to its baseline state is delete the old COW file and create a new (empty) one.

Another key advantage of using COW files is that they allow you to use the same root filesystem image on more than one guest system simultaneously. All you need to do is specify a different COW file each time you bring up a guest kernel. In fact, you can use both the same image file and the same kernel for multiple guests. As you can guess, we're going to use a COW file in our example scenarios.

## Getting a Root Filesystem Image

The procedure for building your own root filesystem image boils down to this:



Imagination. Innovation. Integration.

# Building Beyond Expectations



## AMAX - Your turnkey computing solution provider!

- Design Prototyping
- Production
- Fulfillment
- Branding
- Testing
- High Temp Burn-In



### Quad Core Servers

1U High Density Server - Intel® Xeon® Processor based server with only 17" in depth!  
 2U/5U RAID Servers - Advanced multi-core Intel® Xeon® Processor based server solutions with SATA or SAS drives and up to 18TB of hot swap SATA and 7.2TB hot swap SAS storage capacity.



Supports up to 84 Intel® server blades and 168 energy efficient Dual-Core Intel® Xeon® Processors in one rack!

### Expand Your HPC Flexibility

Enjoy great dual-core and quad-core performance plus the flexibility to run 32-bit and 64-bit applications with Dual-Core Intel® Xeon® Processors and Quad-Core Intel® Xeon® Processors in your **AMAX HPC Servers.**



Xeon® inside™

Dual-core. Do more.



**AMAX Global Services Also Available!**

**Toll Free: 1.866.626.2948 ext.1206**

**WWW.AMAX.COM**

## Using a filesystem that takes the form of a single file has two important ramifications for User-Mode Linux: first, it helps keep your guest system relatively compact and portable; second, it makes change control as simple as tracking changes to a single file, via the COW file method.

1. Create an empty filesystem image file and mount it to some directory.
2. Install Linux into that directory.

Sounds simple, right? On Debian and SUSE it is—sort of. On other distributions, it's much less so. Regardless, I'm going to save a more-detailed discussion of that process for my next column, in which I'll cover what I consider to be advanced User-Mode Linux topics and techniques. In the interests of getting you up and running with User-Mode Linux in a gratifyingly quick manner, for now I recommend you download a prebuilt image.

My favorite source of these is Nagafix Ltd.'s "UML Resources" page (see the on-line Resources) from whence you can download root filesystem images for not only Debian guests, but also Gentoo, Slackware, Fedora, Ubuntu and others. Nagafix makes a reasonable effort to keep these images up to date with security patches, which is a nice touch.

In addition, Nagafix provides an MD5 and SHA hash of each image file it provides. You may miss them if you click directly on the x86 and AMD64 links on the page cited above; instead, use the OS-name links, each of which leads to a page containing links not only to images but also to build logs and hashes, plus handy tips on how to update the images yourself.

I obtained my Debian 3.1 image by navigating to [uml.nagafix.co.uk](http://uml.nagafix.co.uk), clicking on Debian 3.1, and then clicking on the root\_fs and MD5 links to download the files `Debian-3.1-x86-root_fs.bz2` and `Debian-3.1-x86-root_fs.bz2.md5`, respectively. After my downloads were complete (the filesystem image itself is 169MB!), I verified the MD5 signature from within a terminal window with the command:

```
md5sum -c ./Debian-3.1-x86-root_fs.bz2.md5
```

And, now we're ready to boot our virtual guest for the first time. We've got a guest kernel named `uml-guestkernel-2.6.17.3`

(from my previous column's example) and a root filesystem image named `Debian-3.1-x86-root_fs.bz2`. You should already be logged in to a terminal session as a nonroot user. Uncompress the filesystem image with the command:

```
bunzip2 ./Debian-3.1-x86-root_fs.bz2
```

Next, just as a sanity check, try booting your guest system:

```
umluser@host:~> ./uml-guestkernel-2.6.17.3
->ubd0=testcow,Debian-3.1-x86-root_fs root=/dev/ubda
```

If all is well, you should see some User-Mode Linux messages, followed by a longer string of Linux kernel startup messages, ending with a login prompt. Log in as `root`—you won't be prompted for a password. Feel free to poke around a bit; you won't hurt anything that can't be fixed later by starting with a fresh COW file.

To see a list of installed packages, enter the command `dpkg -l |less`. You may be surprised by how few Debian packages are present. Don't worry; you'll be able to install additional packets with `apt-get`, just like on a "real" Debian system. When you're done with your initial exploration, issue the command `halt` to shut down your guest system cleanly. We've got some things to do before your guest system can do any serious work—first and foremost is configuring networking.

### Using Bridged Networking with User-Mode Linux

There are a variety of ways to network UML guests, all of which are described in Rusty Russell's User-Mode Linux HOWTO (see Resources). The best option for using UML guests as network servers is bridging, in which your host system acts like an Ethernet bridge between itself, the UML guests running on it and the outside world.

In a nutshell, the procedure is this:

1. Configure your host's TCP/IP stack as a virtual bridge,

## When in Doubt, Roll Your Own Image

Even if you use a root filesystem image from a trusted source and verify its integrity via an MD5, SHA or GPG hash/signature, the fact is, if you're truly worried about security (we are, aren't we?), you're much better off building your own filesystem image than using someone else's.

I'm indulging in just a little laziness and instant gratification by using a prebuilt image in this article, which I think is justifiable in the larger aim of encouraging UML experimentation and adoption. Just be sure to check your image's hash/signature, and the first time you mount it in UML, run `apt-get dist-upgrade` (or YaST Online Update, yum or whatever update mechanism your guest's distro supports).

Next time, I'll discuss the filesystem image build process in more depth, as well as how to use iptables both on your host and on your guest OSes to add another layer of protection to your virtual machines.

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/](http://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/](http://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



Listing 1.

## Setting Up Bridged Networking

```

root@host# bash -c 'echo 1 > /proc/sys/net/ipv4/ip_forward'
root@host# apt-get install bridge-utils uml-utilities
root@host# ifconfig eth0 0.0.0.0 promisc up
root@host# brctl addbr uml-bridge
root@host# brctl setfd uml-bridge 0
root@host# brctl sethello uml-bridge 0
root@host# brctl stp uml-bridge off
root@host# ifconfig uml-bridge 192.168.250.250 netmask 255.255.255.0 up
root@host# brctl addif uml-bridge eth0
root@host# tunctl -u umluser -t uml-conn0
root@host# chgrp uml-net /dev/net/tun
root@host# chmod 660 /dev/net/tun
root@host# ifconfig uml-conn0 0.0.0.0 promisc up
root@host# brctl addif uml-bridge uml-conn0

```

and then define your “real” network interface as the first “port” on that bridge.

2. For each guest system you intend to run, create a local tunnel interface and define it as another port on the bridge.
3. When you start a guest system, define its virtual Ethernet interface (eth0) to be the tunnel interface you created in the previous step.

Listing 1 shows the precise series of commands this translates to, adapted from David Cannings’ useful article “Networking UML Using Bridging”. All these commands must be executed as root.

The first command enables IP forwarding on your host. Although, technically, bridging happens at a lower level than IP forwarding, they amount to the same thing from the kernel’s perspective. Accordingly, if you have a local iptables policy on your host, you’ll need to add rules to the FORWARD table to enable traffic to and from the tunnel interfaces you attach to the host’s bridge.

The second command (apt-get install...), obviously, installs the Debian packages bridge-utils and uml-utilities. bridge-utils provides the brctl command, and uml-utilities provides the tunctl command. For these commands to work, your host kernel needs to have been compiled with 802.1d Ethernet bridging, IP tunneling, Bridged IP/ARP packet filtering and Universal TUN/TAP device driver support.

The third command in Listing 1 (ifconfig eth0...) may seem a bit scary. It resets your host’s Ethernet interface to a (temporarily) IP-free state. Be prepared for an interruption in local network functionality after you execute this command.

The subsequent six commands, however, will restore it by defining a new virtual bridge device (called uml-bridge), configuring it, assigning your host’s IP address to it (192.168.250.250 in this example), and attaching eth0 to it as a virtual bridge port. If the IP address of eth0 on your host was 10.1.1.10 before you reset it to 0.0.0.0, after issuing the first four brctl commands you would use ifconfig uml-bridge 10.1.1.10 netmask 255.255.255.0 up. At this point, your host should be

able to interact with the outside world in exactly the same way as it did before (unless of course your local iptables policy doesn’t have appropriate FORWARD rules yet).

All right, our host system is now a bridge. All that remains is to attach a tunnel port to it. You should repeat the remaining steps in Listing 1 (starting with tunctl -u...) for each guest system you intend to run.

In the tunctl -u... command, umluser is the name of the unprivileged account you intend to use when executing guest kernels, and uml-conn0 is the name of the new tunnel interface you’re creating.

In the subsequent chgrp and chmod commands, we’re changing the permissions of the virtual tunnel device, always /dev/net/tun, to be readable and writable by our unprivileged account. In this example, therefore, the account umluser belongs to the group uml-net. (On my real-life test system, I instead used the group wheel, which my unprivileged account mick belongs to.)

After setting the new tunnel interface’s IP address to 0.0.0.0 (just like we did with eth0), we define it as another port on the local bridge with that last brctl command.

That’s it! Now when we start the guest system, we add the option eth0=tuntap, uml-conn0 to our kernel command line, which tells the kernel to use the tunnel interface uml-conn0 as its virtual eth0. Our complete example command line, which unlike Listing 1, should be run by a nonprivileged user rather than root, looks like this:

```

umluser@host$ ./debkern ubd0=debcow,debroot
➤root=/dev/ubda eth0=tuntap,uml-conn0

```

After the virtual machine starts, you can assign an IP address to (virtual) eth0 via ifconfig, define a default route via route add... (using the same gateway IP that your host system uses), set DNS lookup information in /etc/resolv.conf, and, in short, configure it in precisely the same way that you’d configure a real Debian system.

Once your virtual machine is successfully communicating with your local LAN and beyond, you should immediately configure apt-get and use it to install the latest Debian patches on your virtual guest. You’ll need apt-get working anyhow to install the network software you’ve just gone to all the trouble of building this virtual machine to run. In the case of our example virtual DNS server, these would probably be the Debian packages bind9 and maybe also bind9-doc. Remember, all of these changes will be made to your COW file, so be sure to specify the same COW file on subsequent startups (or merge it into your image via the uml\_moo command).

Next time, we’ll wrap up this series by discussing additional security controls you can use on your guest systems, a nifty COW file trick or two and, of course, how to create a custom root filesystem image. Until then, be safe! ■


**Resources for this article:** [www.linuxjournal.com/article/9385](http://www.linuxjournal.com/article/9385).

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



# 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*





JON "MADDOG" HALL

# Events for Suits

Maddog continues his framework for a good conference by considering the suits.

**The planning** for the conference at the Universidade Federal de Santa Catarina in Florianópolis, Brazil was proceeding well. Several tracks had been selected for the technical subjects, and the planning committee had put out a call for papers and selected several "invited talks" from speakers on topics they knew would be of interest to everyone—at least to all the techies.

But today, the organizers wanted to plan some conference topics for business people—managers who might not understand free software from a technical perspective and who would be bored by sessions on the brilliance of the emacs text editor.

"What about the *suits*?", asked JR, "What types of things should we do for them?"

I told JR that it is hard to get business people for even one day, and that you have to develop a special program for them. Also, their interests do not lie in technical subjects, but in making and saving money. Often their interests also relate to better

products or customer service that can come from the careful application of free software.

We decided to set up a short four-hour conference for the business people, starting with a breakfast sponsored by a few computer vendors. The sponsorship would pay for the room, food and travel for some of the speakers.

"First, we will discuss briefly what free software is, and make it clear that the real value to the software is the freedom to change it to meet your needs", I said. "Some managers think that low cost is the only value."

We also decided to ask a local computer magazine to send one of its writers to discuss subjects such as "where to use free software in the enterprise" and "how to migrate and interoperate using free software". We knew this writer would be fair to free software and would tell customers the truth about how it would fit in to their environment.

"Next, we need case studies", I said. The best way of convincing business people that something will work is to



## 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](http://www.amltd.com) or call 1.800.648.4452 for a real live person.



M7100 Wireless Family

**AML**  
NEVER COMPLICATED™



show them another similar business person making money with free software. This makes the attendees see success, and later they can become your best case studies for future events, after they have been successful in their own businesses.

"How do we find these case studies?", Carlos asked. I answered, "You can go to the Web sites of your sponsoring companies or of local magazines and see if they have any articles about companies similar to the ones that you want to invite to your event. Often the sponsoring companies would be happy to work with their customers to get them to come to your event, and perhaps they will even sponsor the customer's travel to speak at your event."

Other items where business people want clarification are licensing, where to get support, where to get training and other business issues associated with using free software.

After the meeting is over, the business people can talk to the vendors at the vendor exhibit, so the breakfast room would have to be near the main event.

A lot of conferences do not like to have vendor exhibits, but I like having a small vendor exhibit area just to allow attendees to see "the latest and greatest" of the vendors' wares. It is recommended, however, to tell the vendors that you want tabletop displays, small displays that do not take up much room or resources, and that they should mirror the themes of the conference. If the conference deals with multimedia, you might invite vendors who make sound cards, solid-state music players, midi instruments and so on to your event—particularly if these work with free software. If your theme is rapid development, you might invite vendors of compiler suites, test harnesses and so on.

You also should recommend that vendors send some technical people who can answer technical questions, as well as marketing people.

Do not forget to invite the .org groups. These are often the most popular exhibits—a lot of the .org people are doing some really innovative and fun things. Also remember that .orgs usually have even less money than small start-up companies, so often you have to donate the booth to them or sell it to them at a real discount. And, any money you can save the vendors on items such as electricity and Internet support, which is typically very expensive in large venues, will be appreciated twice over by the .org people.

"What about advertising?", asked Dennis.

Although advertising is key, so is timing. These days, the Web is used to allow last-minute changes to programs, accommodations, travel tips and other things, but unless your Web site motivates attendees to come on the first viewing, you may never get them to come back for a second viewing. So, you need to make sure that enough information is available the first time potential attendees go to your site to make them register, and then update it with small items and changes as necessary.

Things necessary on the first showing of the Web site are location, time, themes for the event, main speakers (and hopefully the main speakers' topics, abstracts and bios), and fees (if any) to attend. The more speakers you have lined up by the time you take your Web site live, the better the Web site is for your event. A Web site with a lot of blank spaces does not inspire people to come to an event.

Although you should not advertise your site too early, you also should not advertise too late, as people make plans and may not be able to attend your event simply due to conflicting arrangements. With earlier warning, they might be able to reschedule the conflicting event or have enough time to talk their employers into sending them to the conference.

Once the Web site is ready, look for places to get low-cost or free advertising. Most Linux and PC magazines and on-line portals have event calendars. Most would be happy to include your event in those calendars. Local and public radio shows also have community calendars where they announce local events for free. University bulletin boards, library calendars and local newspapers are also good places to place small advertisements.

Entrance fees are always a touchy subject. A lot of free software people want everything to be free, not realizing that floor space, custodial

care, security guards, insurance needs, electricity and Internet usage cost money. Some events charge very little to the attendees and get all of their money from sponsors and vendor sales. Often these low-cost events do not supply food to the attendees and instead have some type of meal plan available for a small fee or suggest that people eat outside of the event at a restaurant of their own choosing.

I have seen some free events, such as LinuxTAG in Germany, put together a small bag of goodies, such as donated CD-ROM collections, T-shirts and other donations from vendors, which are then sold to the attendees to raise money. And, some people either pass the hat for donations or raffle off items, such as a T-shirt signed by all the speakers. One time such a T-shirt brought several thousand dollars for the organizers to help cover costs.

Finally, have fun with your event. Putting together a one- or two-day conference should not be a person- and relationship-killing proposition. By planning ahead, you should be able to take the time to plan events without burning out anyone. ■

---


Jon "maddog" Hall is the Executive Director of Linux International ([www.li.org](http://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.

---

# Linux Laptops

Starting at \$799

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



**R<sup>3</sup>**  
Technologies™  
Working hard to bring you the best of IT.

**309.34.CUBED**  
**shoprcubed.com**



DOC SEARLS

# Greater Goods

How classical economics fails to comprehend free and open-source software development. And, how it's making a whole new world that's bigger and better for everybody.

The problem with classical economics is that it centers its concerns at the commerce level, and specifically around transactions.

At the beach last summer, I caught up with my cousin, Charles Crissman, PhD—a veteran scientist, agricultural economist and Deputy Director General for Research at CIP (better known as the Potato Institute), a large international development organization headquartered in Peru. What surprised and gratified me most was learning from Charles that the results of CIP's research and development are open and accessible. They don't want to see their work benefit one government, or one company, to the exclusion of anybody else, no matter who pays for the work. Agriculturally speaking, they are not in the business of building silos or walled gardens. Instead, they are in the business of helping nature. Literally.

In response, I explained how free software and open-source developers aren't just helping nature, but making it. Their work is creating the core, mantle and crust of a new digital world of code growing within and alongside the physical one. I added that this digital world's geologies are created on NEA principles: Nobody owns it, Everybody can use it, and Anybody can improve it.

"Yes", he said. "You're talking about public goods." The term public goods intrigued me. But there was no connectivity at the beach, and we really weren't there to discuss economics anyway. So, I did that after I got home.

Public goods are non-rivalrous, it turns out. In other words, they are not scarce. Consuming any of them does not reduce the sum available to others. Wikipedia adds:

The term *public good* is often used to refer to goods that are non-excludable as well as non-rival. This means it is not possible to exclude individuals from the good's consumption. Fresh air may be considered a public good as it is not generally possible to prevent people from breathing it. However, technically speaking,

such goods should be called *pure public goods*. These are highly theoretical definitions: in the real world there may be no such thing as an absolutely non-rival or non-excludable good, but economists think that some goods in the real world approximate closely enough for these concepts to be meaningful.

Wikipedia also provides a handy way to distinguish public goods from others that differ in excludability or rivalness (Table 1).

Wikipedia says, "information goods, such as software development, authorship, and invention" fall in the *public good* category. Yet it seems that the purpose of free and open-source development is to produce a *common pool resource*. As Craig Burton has often observed, the idea is to create common infrastructural building material that supports whole industries, rather than just one player in that industry. We do this by making goods that become abundant by being both open and in the public domain.

What we look for is a "because effect", which is what you get when more money is made *because* of something than *with* something. For example, more money is made *because of* the Internet than *with* the Internet. Or, in geological terms, more money is made on top of it than inside of it—by many orders of magnitude. Take all the money cable and phone companies make by selling connectivity and transport, then compare that with all the money made on top of that connectivity and transport—that is, *because of* it. The ratio of the latter to the former is absurdly large.

Yet the Net's carriers (at least in the US) still believe the only Internet business worthy of the label is selling the Net itself. When I talk with folks who work for the carriers, they can barely imagine benefits to their incumbency other than making money every way they can *with* the Net rather than because of it. Worse, they don't want to see their

Table 1.

## CLASSIC DIVISION OF GOODS IN ECONOMY (FROM WIKIPEDIA)

		Excludability	
		YES	NO
Rivalness	YES	<b>Private good:</b> good: e.g., food, clothing, toys, cars, products subject to inhabitants and other useful contents	<b>Common pool resource:</b> e.g., sea, rivers, forests, their edible value-adds between first sources and final customers
	NO	<b>Club good:</b> e.g., bridges, cable TV, private golf courses, controlled access to copyrighted works	<b>Public good:</b> e.g., law enforcement, national defense, fire fighting, public roads, street lighting

users doing anything other than consuming services. To them, the Net is nothing more than a pipe between producers and consumers, and their job is to make money by delivering stuff from one to the other. Why is that? Is it just that they are stuck in their ways? Or is there more to the problem than that?

When I talk with economically savvy folks about the goals and effects of free and open-source software—or of the Net itself—I often hear the terms “external”, “externality” and “externalities”. It is not meant in a dismissive way, but rather a positive one. Abundant free software production and use might be seen as a *network externality*, resulting from the *network effects* caused by cost-free goods that are easily obtained and used—which is fine. But there is a cost to this perspective. As Wikipedia puts it ([en.wikipedia.org/wiki/Externalities](http://en.wikipedia.org/wiki/Externalities)):

An externality is a side effect from one activity that has consequences for another activity but is not reflected in market prices. Externalities can be either positive, when an external benefit is generated, or negative, when an external cost is generated from a market transaction.

An externality occurs when a decision causes costs or benefits to stakeholders other than the person making the decision, often, though not necessarily, from the use of common goods (for example, a decision that results in pollution of the atmosphere would involve an externality). In other words, the decision-maker does not bear all of the costs or reap all of the gains from his or her action.

Note the perspective. The view of what’s external and what’s internal depends on where you stand. And, classical economics stands with transactions between sellers and buyers. The diagram shown in Figure 1 from the same Wikipedia page (on externality) makes the point of view clear.

Most of us view markets, and economic activity generally, through the prism of transaction. Or, to retain the triangular metaphor, from the top of the pyramid—that is, from the side of the firm, the seller, the producer, the few who sell to the many.

This explains to me why, countless times on the Gillmor Gang podcast, I go silent or into a rant against the “vendor sports” commentary by other Gang members. They see my main area of concern—free and open-source development and DIY activity on the customer’s or consumer’s side—as external to the work of large producers.

What most of us don’t see is that most free and open-source software development isn’t in a business at all. It’s busy making the stuff that makes the world that everybody lives in. It is pro-business the same way the core of the Earth or the Pacific Ocean is pro-business. Its tides lift all boats, but it is not especially concerned with what any of those boats are up to.

Still, so far we’ve concerned ourselves only with a few of the many goods economists talk about. Other adjectives modifying goods include *durable*, *non-durable*, *intermediate*, *capital*, *consumer*, *experience*, *merit*, *complement*, *substitute*, *scarce*, *positional* and *free*.

Of all those, the one that best applies to what we’re up to is *free*. Wikipedia explains:

The *free good* is a term used in economics to describe a good that is not scarce. A free good is available in as great a quantity as desired with zero opportunity cost to society. A good that is made available at zero price is not necessarily a free good. For example, a shop might give away its stock in its promotion, but producing these goods would still have required the use of scarce resources, so this would not be a free good in an economic sense.

# Expert Included.



Chris codes a wide variety of applications, and he expects his hardware to keep pace. He developed the Silicon Mechanics website to help customers find and configure the right servers for their needs, then developed a suite of fully integrated tools to support the entire production process, from configuration through delivery. Chris likes rack-mount servers based on the Dual-core AMD Opteron™ processor with AMD Virtualization™ because virtualization allows him to develop and test software across multiple operating systems on a single server. In addition, the integrated memory controller, now supporting DDR2, reduces latency for fast memory reads, yielding quick computational processing for increased performance.

When you partner with Silicon Mechanics, you get more than a powerful AMD solution — you get an expert like Chris.



visit us at [www.siliconmechanics.com](http://www.siliconmechanics.com)  
or call us toll free at 866-352-1173

**When I talk with folks who work for the carriers, they can barely imagine benefits to their incumbency other than making money every way they can *with* the Net rather than *because* of it.**

There are three main types of free goods:

- 1) Resources that are so abundant in nature that there is enough for everyone to have as much as he or she wants. An example of this is the air that we breathe.
- 2) Resources that are jointly produced. Here the free good is produced as a by-product of something more valuable. Waste products from factories and homes, such as discarded packaging, are often free goods (see also dumpster diving).
- 3) Ideas and works that are reproducible at zero cost, or almost zero cost. For example, if someone invents a new device, many people could copy this invention,

with no danger of this “resource” running out. Other examples include computer programs and Web pages.

Not surprisingly, this is consistent with the Free Software Definition ([www.gnu.org/philosophy/free-sw.html](http://www.gnu.org/philosophy/free-sw.html)) and Richard M. Stallman’s original distinction between free speech (a free good) and free beer (a private good, given away).

Public infrastructure is a *because effect* of free software, which is created down at the level of nature—the level where we make the digital world. That level is nicely positioned by the Long Now Foundation in the diagram shown in Figure 2.

Although this diagram was created to show differences in the speed of change in civilization, it also shows dependencies. Culture depends on nature. Governance depends on culture and nature. Infrastructure depends on all three.

The problem with classical economics is that it centers its concerns at the commerce level, and specifically around transactions. More is involved than just transactions, and a lot of it happens down at these other layers.

Common, public and free goods, whether or not they are produced by commercial activity, are external to it. But, significantly, they are external below, on the supportive side. And you can’t completely understand the virtues or natures of those lower-level goods in commercial terms, economic or otherwise—just as the science of mechanics cannot explain physics or chemistry, even as it relies on them.

From the perspective of commerce, it is hard (maybe impossible) to comprehend the supportive (and not merely the external) purposes of free and open-source software—or why they are so deeply supportive of economic activity and value creation. It is hard to see how, by their nature, free and open-source software provide deep and supportive culture, governance and infrastructure for all kinds of commercial activity. Yet this is how, at the deepest level, we are making the digital world.

The big brain-twister is, it only gets larger. That’s because, unlike the physical world—with its fixed dimensions and its portfolio of building materials assembled from the periodic table of elements—the digital world can be improved by anybody ready and able to contribute useful code.

That code isn’t just in the form of programming, either. It’s in the form of text, music, video and other arts that contribute to common understanding. Here is where we are only beginning to develop the culture and governance that will support new social infrastructures, including those of government and business.

Wikipedia is a perfect example. I’ll be curious to see how the entries on economics that served as sources for this column will change as readers of *Linux Journal* (and other instruments of understanding) make corrections and improvements.

In the old pre-digital world, about all we could do was consume and complain. Now we can produce and construct. And that makes a world of difference. ■

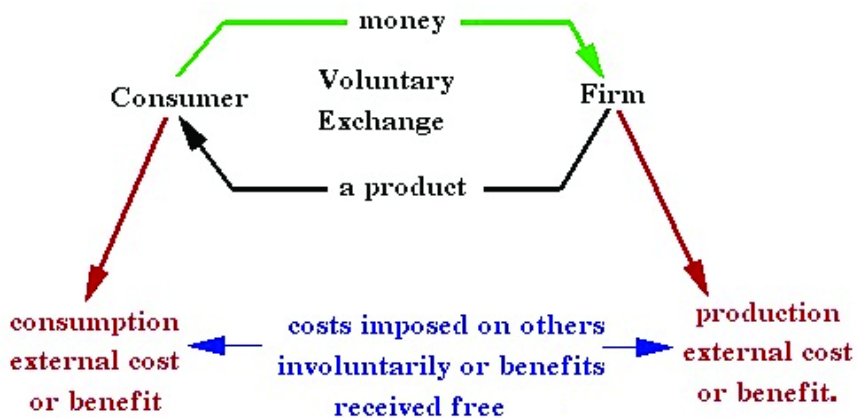


Figure 1. Externality

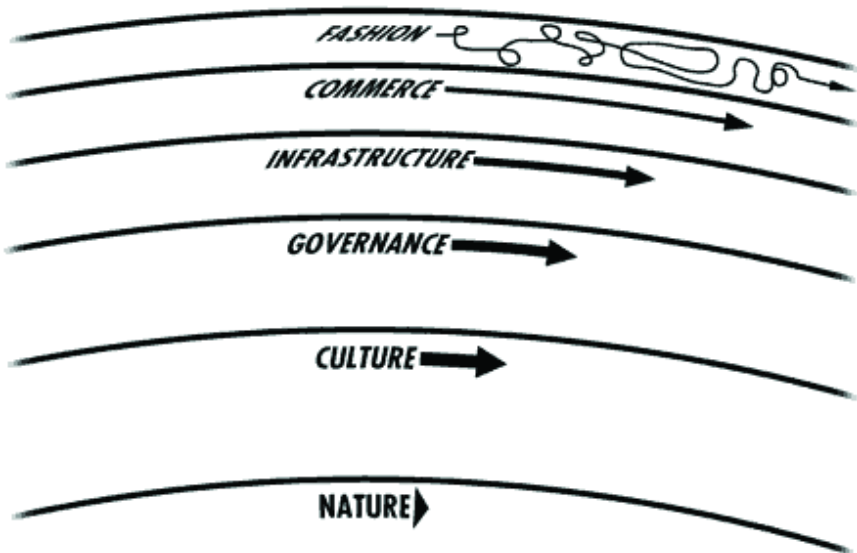


Figure 2. Nature is the level where we make the digital world.

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



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](http://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](http://www.coraid.com)  
for more information.



1.706.548.7200

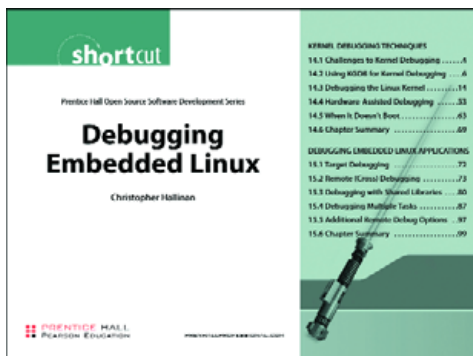
The Linux Storage People

[www.coraid.com](http://www.coraid.com)

## ITTIA DB

ITTIA just released version 1.1 of ITTIA DB, the firm's self-titled, flagship database for deployment in mobile and embedded platforms. ITTIA says that its fully cross-platform database offers developers "fine-grain control over how system resources are used in order to produce efficient mobile and embedded applications...where the limited memory, storage and processing power requirements make software development challenging." This upgraded version boasts an enhanced C API, increased control over storage size for each file type, an improved interface for accessing BLOB data, modified transaction handling for improved tracking of resource-acquisition bugs and other performance and configuration enhancements. ITTIA notes that many customers utilize its product on embedded Linux platforms, for instance, "HVAC controller systems, physical access control devices and consumer electronics". You can get an evaluation copy of ITTIA DB from the company's Web site.

[www.ittia.com](http://www.ittia.com)



## Pearson Technology Group's Digital Short Cuts

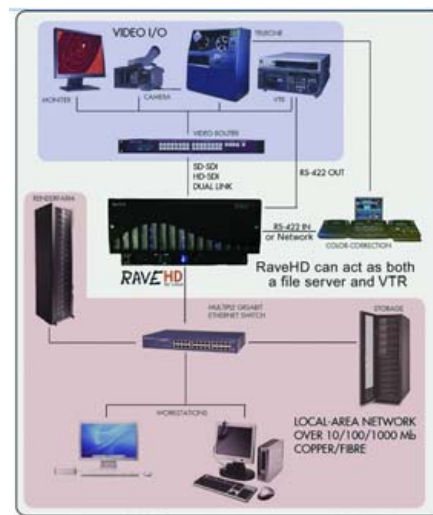
Getting cutting-edge IT information from an author's brain to yours more quickly is the mission of Short Cuts, a new line of digital documents from Pearson Technology Group (PTG). Short Cuts are "concise PDF documents about a cutting-edge technology that shows great promise, or an existing technology that has reached the 'tipping point' and is about to take off", says PTG. The rationale is that when a hot topic comes along, many readers don't want to wait the extra weeks or months needed for the information finally to reach the printed page. Despite the rapid availability, PTG claims that Short Cuts retain the "same level of quality, accuracy, knowledge, and insight" as printed books. The titles span a wide range of IT topics from Pearson's various imprints, including Addison-Wesley Professional, Cisco Press, Exam Cram and Prentice Hall Professional, among others.

[www.informit.com/shortcuts](http://www.informit.com/shortcuts)

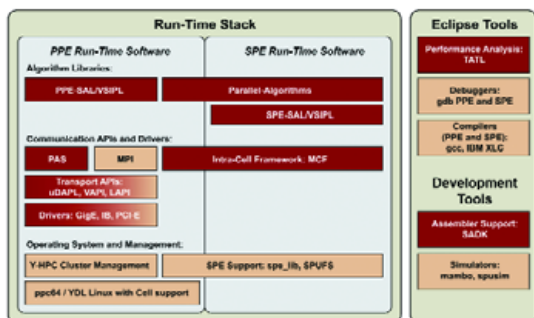
## SpectSoft's RaveHD

Yes, folks, RaveHD is a bit esoteric...but that's what makes it so cool! RaveHD's producer, SpectSoft, recently released a major new upgrade to its non-version-numbered product, which is a combination video transport recorder (VTR) and file server for film production. Utilizing Linux and its own in-house software app, RaveHD stores industry-standard DPX frames and makes them accessible via the network, or it can feed those frames to an onboard I/O board as a video stream. DPX frames allow time-code, audio and other material to be packed into each individual frame. The RaveHD hardware must sustain 300Mbps for a video stream for both ingest and playout. However, the hardware exceeds this by far, making RaveHD an ideal file server to feed these frames into other apps. Other tools support particular work flows in the film industry, "such as VFX for dailies and feature film for ingest on-set", says SpectSoft. RaveHD's latest major features include an auto-router, which "allows the easy routing of any of the SD, HD or Dual Link formats to the various features within the I/O board", as well as a JPEG push that converts any frame to a JPEG and pushes it either to the RaveHD GUI or any browser. Hey, Mom, I know what I want for Christmas!

[www.spectsoft.com](http://www.spectsoft.com)



RaveHD in a standard multi-tier environment can be accessed via RS422 and standard network.



## Mercury Computer Systems' MultiCore Plus SDK

The MultiCore Plus SDK from Mercury Computer Systems, now free from the bonds of beta at version 1.0, is a seamless package of software development tools and libraries that enables its users to exploit the Cell Broadband Engine (BE) and other multicore processors fully. According to Mercury, the SDK "includes a comprehensive programming framework, highly optimized math libraries and a graphical IDE with powerful debug and analysis tools". Furthermore, supported on the open-source Linux distro for the Cell BE processor, the SDK complements components of the IBM SDK. The beta version of the product has been present in applications, such as aerospace and defense, seismic/geologic, semiconductor, life sciences, digital media and national labs. Both Mercury and IBM also offer a range of Cell BE processor-based products.

[www.mc.com](http://www.mc.com)





## AML's M5900 Series Portable Data Terminal

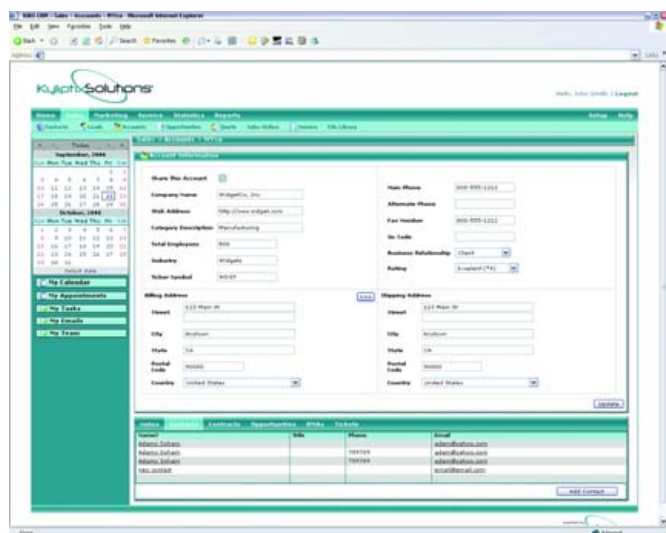

AML has graced this page numerous times with its offerings, and this time around it has a new data-capture device, the M5900, which aims to "supply big-business functionality at a small-business price". AML's target customer is one needing "high performance for everyday, all-day data collection applications, including inventory control, factory-floor management, price verification, shipping/receiving, asset tracking" and so on. Feature-wise, one will find 32MB RAM/16MB Flash ROM memory (with 10MB of user-available non-volatile memory), a 200MHz ARM9 processor, a rechargeable lithium-ion battery (plus backup), backlit LCD display, a 55-key keypad and an SQLite database engine—with an embedded Linux OS running the show, of course. Other options include industrial or general-purpose configurations, as well as four different laser choices.

[www.amltd.com](http://www.amltd.com)

## Joseph Weber and Tom Newberry's IPTV Crash Course (McGraw-Hill)

Getting your TV fix delivered to you via IP is becoming ever more common, and one way to understand that universe better is with Tom Newberry and Joseph M. Weber's new book, *IPTV Crash Course*. This work is an "accessible overview" of IPTV—that is, the convergence of the Internet and digital video technology. Its mission is to "explain the fundamentals of IPTV", as well as "how the business models of service carriers will change" due to the utilization of new technologies. Although much of the tech stuff will be familiar to most of us, the societal and economic impacts that are covered here are likely to tickle both the suit and the geek alike.

[books.mcgraw-hill.com](http://books.mcgraw-hill.com)



## Kyliptix Solutions' KiBS CRM

The KiBS CRM is a Web-enabled, SaaS-based CRM module for small- and medium-sized businesses, offering "integrated sales, marketing, customer service and support" together in one package. It is the first application in the Kyliptix Integrated Business Suite (KiBS), which is targeted at small- and mid-sized businesses. Kyliptix claims that KiBS "is capable of integrating with existing front- and back-office applications", meaning that customers are "no longer forced to engage a system integrator to create problematic patch code to ensure interoperability and communication between the multiple software applications". By working with existing data rather than replicating or porting data to other locations, says Kyliptix, "KiBS eliminates compatibility issues and errors stemming from improper synchronizations". KiBS is built upon a LAMP platform and utilizes an Ajax methodology. Additional modules are forthcoming, according to the company.

[www.kyliptix.com](http://www.kyliptix.com)

Please send information about releases of Linux-related products to James Gray at [newproducts@ssc.com](mailto: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.



# Editors' Choice

# 2006

Excellent competition forced our editors to make some difficult decisions for the 2006 Editors' Choice Awards for software, hardware and services. LJ Staff

**O**utside the open-source products that work on Windows as well as Linux, Windows continues to evolve into a Microsoft-only platform, as Microsoft continues to eliminate its commercial competition. But, competition thrives more than ever on Linux. Pick just about any category of software, and you can find at least two or three excellent candidates, often more. And, the available products are so good, it's difficult to claim that there is a hands-down best, such as the definitive word processor or e-mail client. One's choice often boils down to personal taste.

Here's how we approached the task of selecting winners this year. We asked vendors to nominate the products they released this year. We combined

their nominations with our own choices, and our editors chose the best of the best for each category. If there's a flaw in our nomination process, it is that it is sometimes impossible for our editors to try out every product or service in the list of nominations. This is especially true of things like hosting and colocation services. In cases like these, personal experience had to trump vendor claims. For example, there may be a better hosting service than the one that earned our award, but we can vouch for our choice from personal experience, which carries more weight.

In the end, the process was fun, despite the challenge. We hope some of your favorites captured top honors, or if not, managed to get an honorable mention. So, on with the show.

# Linux Distribution

## Ubuntu 6.06

It was at once the easiest and most difficult decision to pick the distribution for the Editors' Choice. Ubuntu has a long list of features and design decisions to recommend it for our award. It is easy to install; it has a vast repository of software; it is stable and friendly; it protects users from logging in as root by default and much more. One of the most influential factors in our decision was the fact that Ubuntu has captured and held more popular interest than any other distribution almost since its release. Granted, this isn't a people's choice award, but it's not for nothing that Ubuntu is such a popular distribution. Many of us at *Linux Journal* run it, or its KDE-based sister Kubuntu, ourselves.

Nevertheless, the competition is so superb that a proper list of honorable mentions would be uncomfortably long. We should consider ourselves blessed that we have such a marvelous variety from which to choose. Although we can't name every distribution we could consider worthy of the Editors' Choice Award, we can't resist giving honorable mention to a few. Novell's SUSE Linux Enterprise 10 is arguably the strongest comprehensive commercial distribution available. Linspire could be the ultimate desktop-oriented distribution for new users, although Xandros gives it a run for its money. Gentoo is the definitive compile-it-yourself distribution. Debian deserves a long round of applause, especially since many of the most excellent distributions, including Ubuntu/Kubuntu, Linspire, Xandros, MEPIS, Knoppix and many more are based on Debian. rPath uses Fedora as the foundation for its roll-it-yourself distribution—a perfect choice for those who need to produce custom appliance-like distributions. Even Damn Small Linux deserves a mention for being one of the few distributions that still runs well on older hardware.

As difficult a decision as it was, however, we're more than satisfied with our choice of Ubuntu 6.06 for the Editors' Choice of 2006.

[www.ubuntu.com](http://www.ubuntu.com)



# Office Suite

## OpenOffice.org 2.0.3

There's a joke among musicians that Beethoven wrote only three symphonies: the third, the fifth and the ninth. These three eclipse the rest in terms of popularity such that most people are unaware the other symphonies exist. So it is with OpenOffice.org. OpenOffice.org is so popular, it eclipses the competition to the point that many people are unaware there is competition.

For example, Evermore Software's EOffice suite has superior live links and duplicates the Microsoft Office interface almost exactly, but it is not open source, and it isn't marketed aggressively enough such that many people know it exists. KDE's KOffice suite is a powerful suite of productivity applications, but it is often overlooked because it doesn't attempt to mimic Microsoft Office.

OpenOffice.org delivers just the right combination of openness, power and similarity to Microsoft Office that it provides the features and familiarity people want in an office suite without the drawbacks of proprietary document format or proprietary code. It may not always import Microsoft Office files perfectly, but it does so without the crashes that sometimes plague suites like EOffice when importing large, complex Microsoft Office files. Overall, OpenOffice.org has a way to go before it reaches its potential, but it still provides the best combination of features and compatibility, along with the distinct advantage of being an open-source project.

[www.openoffice.org](http://www.openoffice.org)

# Desktop Environment

## KDE 3.5.4

KDE is the desktop with everything. It is friendly, intuitive and simple enough for the casual user who wants to use it as-is, but it also packs nearly unlimited features and configurability for those who want to plumb the depths of its power. For example, click on the default Konqueror button, and it takes you to a default page with links to your home folder, network folders, applications, trash bin and storage media. Click on the home folder link, and you get a simple, intuitive, folder-based file manager.

That would be enough for most people, but power users who want more from Konqueror can open a navigation panel, split windows multiple times, open tabbed panels—there's almost no limit to what you can do. You can use the fish: kio-slave to view and manipulate files on another computer over a secure connection. And, when you're happy with a view into your own filesystem or that of another computer, you can save any combination of URI and window configuration as a profile you can restore instantly.

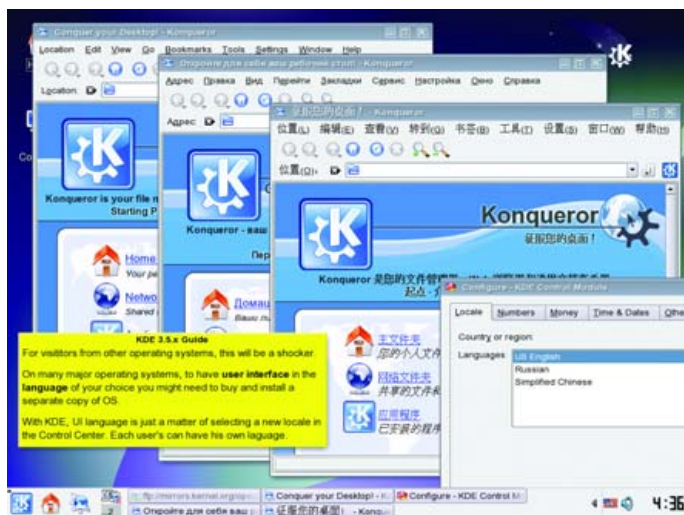
Or, as another example, you can pop an audio CD into your CD drive, and Konqueror opens a window with virtual folders of your songs in MP3, Ogg Vorbis and other formats (depending on which extensions you have installed). Ripping your songs to MP3 format is as simple as copying and pasting the virtual MP3 files to another folder or to your MP3 player.

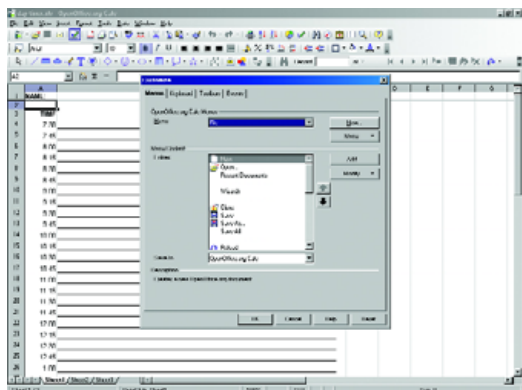
According to research organizations such as Evans Data, KDE is the most popular desktop environment. How does that square with the fact that GNOME is the default desktop of one of the most popular distributions (Ubuntu)? We have no idea. Whether or not Ubuntu users are sticking with GNOME or installing KDE, GNOME certainly deserves an honorable mention on its own merits. GNOME has come a long way in

recent times, and it is particularly appealing in its default Ubuntu configuration.

GNOME was first to integrate the Beagle search engine into the desktop. Beagle is a Mono-based adaptation of the Java-based Lucene search engine. It is capable of indexing files of a wide variety of formats, so you can search through the contents of those files almost instantly. KDE has a search tool called Kerry, which is the equivalent of the GNOME search tool. Although GNOME should get credit for introducing the feature, KDE's power is made more apparent by how KDE easily integrates Beagle into Konqueror as a kio-slave. Put simply, you can type `beagle:ubuntu` in the Konqueror location bar (where you might type a file path or Web URL), and Konqueror taps into the Beagle search index to find all files containing the word ubuntu. All the files found will show up in the Konqueror window as icons and previews.

[www.kde.org](http://www.kde.org)





## Spreadsheet

### OpenOffice.org 2.0.3 Calc

OpenOffice.org Calc makes good on the same formula that has made the entire suite so successful. It is an excellent blend of power and compatibility with Microsoft Office, and it has the added bonus of being based on open source and open document formats. Gnumeric and KSpread deserve honorable mentions, but if you're really serious about doing spreadsheet work, your best bet is with OpenOffice.org Calc. [www.openoffice.org](http://www.openoffice.org)



## Web Browser

### Firefox 1.5.0.6

Is there really any other choice but Firefox? Actually, there are good alternatives. Konqueror is reportedly faster than Firefox. Opera is no slouch in terms of speed and features either. But, after all is said and done, Firefox is the clear winner, and one of the easiest decisions we had to make for Editors' Choice. How do you beat a browser that can please virtually everyone? Choose your favorite theme and add a few extensions, and you can make it look exactly the way you want it to look and do just about anything a browser can do.

Firefox is easy to use, easy to install, speedy, compliant with standards and compatible across different platforms. As we mentioned, it has a huge number of extensions that allow you to customize it for your own tastes and needs. Do you want to synchronize bookmarks from one platform to another? Grab an extension, and you can use commercial bookmark storage or your own server as the central repository for your bookmarks, depending on the extension you choose. Install another extension to manage all your passwords. Install yet another extension to have Firefox check your Gmail or Yahoo accounts. Some favorite geek extensions are Web developer, Book Burro (to feed your used-book-buying habit), Firebug (for modern JavaScript debugging), Flashblock (to stop Flash animation from starting until you want it to), Session Saver (to return to the same set of windows) and Greasemonkey (for client-side JavaScript programs). There is practically no limit to what you can do with Firefox. [www.mozilla.com](http://www.mozilla.com)

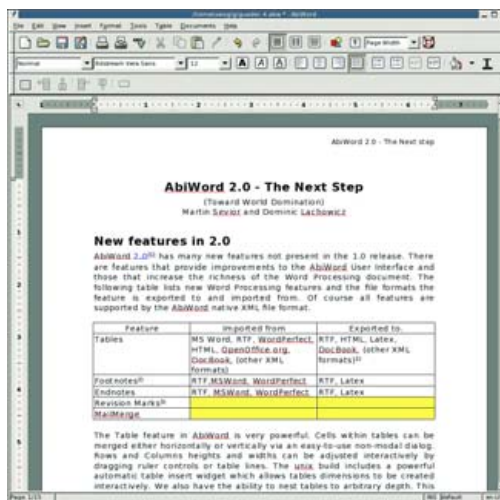
## Word Processor

### AbiWord 2.4.4

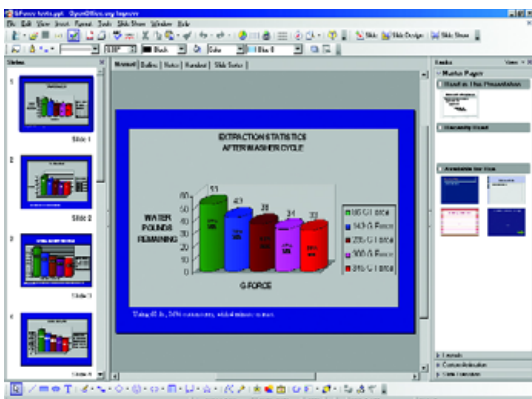
Here is where we break tradition and give the Editors' Choice Award to a productivity application that doesn't appear in the OpenOffice.org suite. Some of us just want to do word processing. We don't use spreadsheets or create presentations, so it isn't important to have a full office suite. We want a word processor that is lean and mean, starts up faster than OpenOffice.org Writer, imports Microsoft Word files adequately and offers all the features we need.

Two word processors fit the bill nicely: KWord and AbiWord. We could justify giving either of these the Editors' Choice Award. We went with AbiWord 2.4.4 primarily because it has a slightly more familiar look and feel for Microsoft Word users, and because it sports a number of very useful plugins. For example, one plugin allows you to place the cursor on a word and run a Google search on that word. Another lets you look up the word in Wikipedia. Still another is supposed to translate selected text via Babel Fish, although that plugin wasn't fully automated in our experiment. Still other plugins add the ability to read and write various document formats, including OpenOffice.org Writer files and Microsoft Word.

AbiWord has all of what most people will need in a word processor and then some, without the bloat and long load times of OpenOffice.org Writer. [www.abisource.com](http://www.abisource.com)



## Presentation Software



### OpenOffice.org 2.0.3 Impress


We came back to OpenOffice.org for our choice of presentation software. As with the spreadsheet and the entire suite, it offers that optimal balance of features, power and familiarity for those who want to migrate from Microsoft Office. And, of course, it has the oh-so-important benefit of being open source and supporting an open document format. KPresenter deserves an honorable mention, as does the presentation module in EIOffice. [www.openoffice.com](http://www.openoffice.com)

# Database

## PostgreSQL 8.1.4

Time to put on our flame-retardant suits. How could we pick any database other than MySQL? MySQL is the M in the LAMP stack. Much of the Web practically runs on MySQL. But, we continue to be most impressed by the open-source PostgreSQL. It handles everything we throw at it and just keeps working, flawlessly. It's almost invisible from an administrative perspective. It handles huge quantities of data, and it has all of the goodies that we expect in a relational database (such as referential integrity, column-level constraints and checks, server-side functions, subselects and unions). The original 8.1 release, which came out in November 2005, included a number of new features, such as two-phase commits. We can't recommend PostgreSQL highly enough.

Having said all that, MySQL certainly deserves an honorable mention at the very least. It is a staple and deservedly so.  
[www.postgresql.org](http://www.postgresql.org)



# Security Tool

## Novell AppArmor

AppArmor strikes a reasonable balance between the complexity and power of SELinux and Linux's default "winner/root takes all" security model. With its wizard-based setup tools (integrated into SUSE's YaST system administration GUI), AppArmor makes it easy even for nonsecurity geeks to strengthen their mission-critical applications with kernel-level mandatory access controls.

AppArmor is included in recent versions of SUSE Linux, including the free OpenSUSE distribution. Although at present AppArmor runs only on SUSE, Novell has released AppArmor's source code (which it acquired from Immunix) licensed under the GPL. Efforts are underway to port it to Ubuntu (and therefore also Debian); other ports should follow.

PacketFence deserves a mention here too. Finally, we have a well-structured tool that combines the power of many open-source components to do network policy enforcement.  
[www.novell.com](http://www.novell.com)



# Mail Client

## Thunderbird 1.5.0.5

It may be a no-brainer to pick Firefox for Editors' Choice, but the decision to elect its sister mail program, Thunderbird, was far more difficult. There's no lack of good e-mail clients for Linux. Evolution and Kontact are not only excellent e-mail clients, they include calendars and other nice features—not that features make the e-mail client. Heck, some of us at *Linux Journal* still think the character-based Mutt is the bee's knees.

We ended up choosing Thunderbird for some of the same reasons we picked Firefox—extensibility. You may be satisfied with Thunderbird right "out of the box". But are you frustrated when you get an e-mail with a URL that is broken into several lines so that you can't just click on it to bring up the Web page? Install the URL Link Thunderbird extension—problem solved. Is the default spam filter for Thunderbird failing to catch all your spam? Install the Spamato4Thunderbird extension—problem solved. Although there aren't as many extensions for Thunderbird as there are for Firefox, and the best expansion is yet to come (the Lightning calendar extension is still in the works). There's enough flexibility in what you can do with Thunderbird to make it suit almost any taste.


Nevertheless, we gladly award honorable mentions to Evolution, Kontact and, yes, even Mutt.  
[www.mozilla.com](http://www.mozilla.com)

# Language

## Ruby 1.8.5

Not since Python has any language captured the imagination of so many eager programmers. Ruby is an object-oriented scripting language that is natural, easy to work with and, well, fun. Ruby on Rails expanded the awareness of Ruby as a language, and now Sun has blessed JRuby (Ruby implemented in Java) by hiring two JRuby developers to work on it full-time. The bottom line is: Ruby is going places, and it is likely to be headed for explosive popularity. People who want in on the fun should grab a copy and start learning it, lest they get left behind when the revolution comes.

Some of our editors would stage a revolt if we didn't give honorable mentions to Objective-C, Perl and Python.  
[www.ruby-lang.org](http://www.ruby-lang.org)



# Game/Entertainment Software

## Quake 4

This AAA (top-tier) game title offers a native Linux client with no compromises from the Windows version, so Linux users aren't getting a second-class product. id Software has released Linux versions for all versions of *Quake* and later versions of *Doom*, which will hopefully catch the attention of other major game publishers.

TransGaming Software gets an honorable mention for its work in allowing Linux users to play popular, non-Linux AAA titles, such as *World of Warcraft* on Linux without having to dual boot.  
[www.quake4game.com](http://www.quake4game.com)



## Web Server

### Apache 2.2

Are there really any other serious contenders for Editors' Choice of Web server for Linux systems? There are other open-source alternatives, such as the AOL server, but Apache still enjoys the most language and module support. It may be the extensions and add-ons that make Apache interesting as a Web development platform, but as Apache is the de facto standard engine of choice, it would be hard to justify giving any other Web server the Editors' Choice Award. Lighttpd deserves an honorable mention. It is becoming popular for its good FCGI support, which is used in Ruby on Rails. [httpd.apache.org](http://httpd.apache.org)



## Communication Tool

### Asterisk 1.2.12

Asterisk is an open-source, complete Private Branch Exchange (PBX) with a list of features that won't quit. It is currently maintained by the Debian VoIP Team and sponsored by hardware vendor Digium. Digium makes hardware that works with Asterisk, but Asterisk works with hardware other than Digium's product line. Asterisk is a no-brainer for Editors' Choice if there ever was one. Features out the wazoo, completely open source, free to use—what more could one hope for in a VoIP solution? [www.asterisk.org](http://www.asterisk.org)



## Web Application Framework

### Ruby on Rails 1.1.6

Not only has Ruby on Rails skyrocketed in its acceptance during the last few years, but people who use it generally fall head over heels in love with it. Some developers say they look at old Web applications they wrote using other frameworks and almost start crying when they discover that Rails could have eliminated 50–70% of the code that went into those projects. [www.rubyonrails.org](http://www.rubyonrails.org)



## Software Development Tool

### Eclipse 3.2


Eclipse is a Java-based extensible integrated development environment (IDE). According to several Evans Data Corporation surveys, it is the most popular development environment among professional Linux developers. To say that Eclipse is extensible is almost an understatement. There are plugins to make Eclipse do just about everything except groom your dog (although we hear that plugin is in the works).


Another honorable mention goes to VMware Workstation 5.5. Virtualization has revolutionized the way we test and provision operating systems, and VMware is still the most mature, versatile and easy-to-use cross-platform virtualization environment. VMware has a long history of working as well or better on Linux hosts as on Windows. And, nowadays it's free too. VMware has made VMware Server (though not VMware Workstation) a free download. [www.eclipse.org](http://www.eclipse.org)

## Development Book

### Ajax Design Patterns by Michael Mahemoff

*Ajax Design Patterns*, published by O'Reilly, assumes that you have a good idea of how HTTP, HTML, the DOM and CSS work (although it does help you brush up as necessary), and it shows you how to combine the basics into sophisticated applications. You can almost think of it as an Ajax cookbook, but with the underlying theory and advice that you need to make interesting applications. [www.oreilly.com/catalog/ajaxdp](http://www.oreilly.com/catalog/ajaxdp)






## Notebook

### Lenovo ThinkPad T Series

The Lenovo ThinkPad models from the T series are relatively inexpensive, durably built, and the driver support in Linux is very good. Wireless and wired network support, video and sound work well with most recent distributions out there. These laptops run solidly for years and perform very well.  
[www.pc.ibm.com/us/notebooks/thinkpad/t-series/index.html](http://www.pc.ibm.com/us/notebooks/thinkpad/t-series/index.html)



## End-User or Non-Technical Book

### **Beginning Ubuntu Linux: From Novice to Professional** by Keir Thomas

What better complement to the Editors' Choice for Linux distribution than a book on how to use that distribution? This book by Keir Thomas, published by Apress, is such a handy resource that we published a sample chapter in our October 2006 issue.

[www.apress.com/book/bookDisplay.html?bID=10086](http://www.apress.com/book/bookDisplay.html?bID=10086)

## Management or Admin Software

### Mantis Bug Tracking System 1.0.5

When thinking of management or administration software, bug tracking might not immediately pop into mind. But the Mantis Bug Tracking System can be invaluable in a corporate environment where much of the company relies on in-house development to keep the business afloat. Mantis is a PHP Web-based tool that is easy to install, intuitive to use and handles multiple projects.

[www.mantisbugtracker.com](http://www.mantisbugtracker.com)

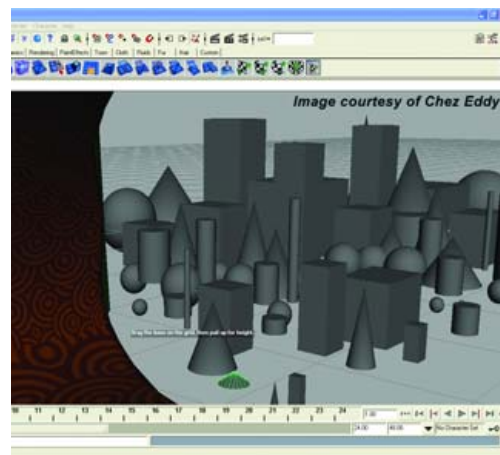


## Graphics Software

### Autodesk Maya 8

Autodesk Maya is an integrated 3-D modeling, animation and rendering solution. It rendered the animation and special effects for movies such as *The Chronicles of Narnia*. Version 8 is the first full release of Maya that runs on 64-bit Linux, a milestone that makes the software even more compelling. If Maya 8 is out of reach of your budget and/or ambitions, Toon Boom Animation, Inc. ([www.toonboom.com](http://www.toonboom.com)) sells a wide variety of 2-D and 3-D animation software, with packages for home users to studio professionals. The Toon Boom products are all available for Linux. Any of these could have been our second choice.

[usa.autodesk.com](http://usa.autodesk.com)



## Hosting or Colocation

### Johncompanies

Our Editor has used Johncompanies for years and testifies that they're wonderful. No technical question is too hard for them. It's a bit creepy that you don't know much about the company other than the name of the head honcho (John) and the Linux technical support person (Dave). Even the sales staff goes by JC Sales. And, instead of a Web or e-mail ticketing system, they simply answer e-mail, which seems like it shouldn't work. But in the time that our editor has been using Johncompanies, it has been competent, friendly and helpful, surpassing other hosting services by a very large degree.

[www.johncompanies.com](http://www.johncompanies.com)



## Mobile Device

### Funambol

Funambol isn't actually a mobile device, but we chose to give it the Editors' Choice if for no other reason than to avoid plugging the Nokia 770 yet again. Funambol is an open-source SyncML server that acts as a middle-ware between groupware servers and mobile devices. It supports the most popular PDAs and commodity mobile phones. It's great, and the community is finally coming up with a solution that rivals the best commercial competition. Check out the Web site for more information.

[www.funambol.com](http://www.funambol.com)

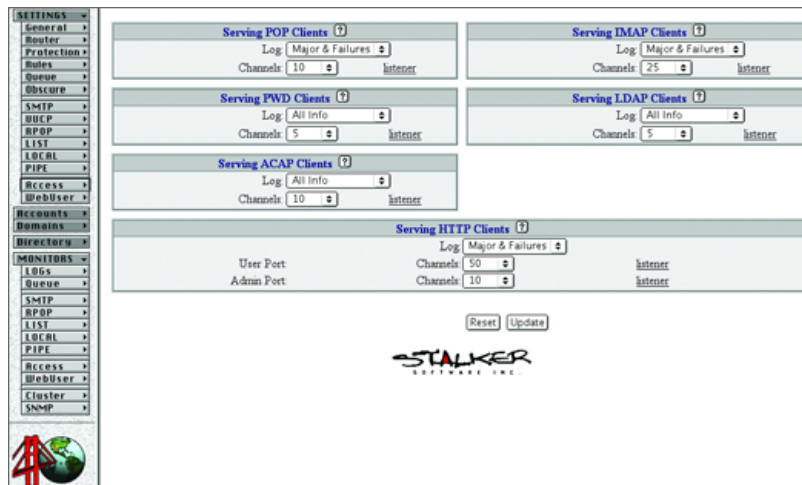
## Software Library or Module

### Yahoo UI (YUI) Library

Under normal circumstances, Qt 4 would be a shoe-in for Editors' Choice in this category. Considering how important Ajax has become to development, we chose the rich library released to open source by Yahoo. It is a comprehensive library of components, utilities, controls and CSS resources for the Ajax and Web services developer.

The Google Web Toolkit ([code.google.com/webtoolkit](http://code.google.com/webtoolkit)) was a close second. Google released a lot of its resources under open source, although a few goodies are still missing. For example, the hooks are there to create something like the drag-and-drop gadgets you can assemble on your personal Google page, but we suspect Google has some unreleased code to make this much easier than what you have to do to make it work with the currently released GWT.

Honorable mention goes to Prototype ([prototype.conio.net](http://prototype.conio.net)), a JavaScript library that makes it easy and fun to work with JavaScript. Prototype has become famous in part because of its inclusion in Ruby on Rails. But, you can use Prototype without Rails, and Prototype itself is the basis for some higher-level projects and libraries, such as Scriptaculous. If you work with JavaScript, you should check out Prototype. [developer.yahoo.com/yui](http://developer.yahoo.com/yui)




## POP/IMAP/SMTP E-Mail Server

### CommuniGate Pro 5.1

CommuniGate Pro 5.1 is a comprehensive Internet communications system that encompasses IMAP, POP, SMTP, groupware and even includes a VoIP PBX. CommuniGate Pro always has been one of the easiest servers to manage, and that ease of use has been extended to its new VoIP capabilities. It is a cinch to create outgoing messages, voice menu systems, call conferencing, caller-ID blocking and much more. It is outrageously simple to set up CommuniGate Pro clusters, making it one of easiest solutions for situations where scalability is important. Of course, it supports the gamut of e-mail features, including LDAP directories, Web mail, hooks into antivirus software and spam blockers, and an easily configurable set of filters. If you're allergic to proprietary commercial software, you'll want to avoid this one, but you'll have to put in a lot of time and effort to duplicate with open source what you can get so easily with CommuniGate Pro.

The Gordano Messaging Suite ([www.gordano.com](http://www.gordano.com)) is a commercial Exchange replacement alternative that features instant messaging, collaboration, mobile gateway and archive/recovery. If you want an open-source solution, Open-Xchange Server 5 ([www.open-xchange.com](http://www.open-xchange.com)) deserves the honorable mention. Open-Xchange server is a terrific open-source drop-in replacement for Microsoft Exchange. It's a classy product for what it delivers. Although it is open source, it is not free. In fact, a year's subscription to the maintenance portal for 25 users, at \$1,095 US, is more expensive than the more feature-rich and scalable 25-user CommuniGate Pro server, which sells for \$699 US. [www.communiGate.com](http://www.communiGate.com)



## System Administration Book

### LPI Linux Certification in a Nutshell, Second Edition, by Steven Pritchard, Bruno Gomes Pessanha, Nicolai Langfeldt, Jeffrey Dean and James Stanger


This O'Reilly book can help you pass your LPI exams or just assist your progress toward being a better Linux system administrator. We'd love to give honorable mention to two other O'Reilly books: *Linux Server Security*, Second Edition, by our own Michael D. Bauer, and *Linux Server Hacks*, Volume Two, by William von Hagen and Brian K. Jones, but both books were released in 2005. [www.oreilly.com/catalog/lpicertnut2](http://www.oreilly.com/catalog/lpicertnut2)

### Contributors to the Editors' Choice Awards

Nicholas Petreley, Dee-Ann LeBlanc, Paul E. McKenney, Michael D. Bauer, Ludovic Marcotte, Mark Brownstein and James Gray.



# ARE YOU EQUIPPED TO DEBUG IN A MULTICORE ENVIRONMENT?



Characters and Images ©2006 Brad Fitzpatrick, ActiveEdge. All Rights Reserved.

*Multicore chips present unique debugging challenges...are you prepared?*

TotalView can help you find even the most elusive bugs and memory leaks – *fast* – in multi-threaded, high performance, distributed or cluster computing environments. Available on UNIX, Linux and Mac OS X.

**Get your TotalView evaluation copy now at [www.etnus.com](http://www.etnus.com).**

**TotalView...THE Debugger for the Multicore Age.**



## LyX and Lulu

Use LyX to create stellar on-line books for the Lulu publishing service. DONALD EMMACK

**Writers are everywhere.** They can be in school, in business or trying to make a living printing text on the page. The Internet abounds with various Microsoft Windows tools to aid in writing books, transcripts and other media. Many of these sites and programs still rely on traditional word-processing programs for output.

LyX is different. It's a typesetting tool designed on LaTeX. In short, LyX makes your printed documents look more like what comes from a professional publishing company. Lulu.com is a fast-growing Web site where you can publish that book you've been meaning to write for the last ten years.

Together, LyX and Lulu make a great pair. Although they can't fix your poor writing habits, they will make your final publication look professionally printed and bound.

This is part one of a two-part series. In this first article, I explain some of the striking benefits of LyX and how to get your final document into the Lulu.com Web site. The next article will focus on using Pixel to create your final book cover for the publication.

### What's LaTeX?

LaTeX is a typesetting system, not a word processor. Word processors fit nicely in the business world, because they give command of fancy document layout to the end user. They also have other tools you expect, such as spell checkers or an automated thesaurus.

LaTeX did not impress me at first. Its raw form is ugly and difficult to understand. Just looking at the text, you cannot realize what it will look like in final printed form. Consider the following text example from the LaTeX Web site:

```
\documentclass{article}
\title{Cartesian closed categories and the price of eggs}
\author{Jane Doe}
\date{September 1994}
\begin{document}
  \maketitle
  Hello world!
\end{document}
```

This is what you need to type into a text editor for LaTeX to render a graphical output. But what you get after using LaTeX is:

```
Cartesian closed categories and the price of eggs
Jane Doe
September 1994

Hello world!
```

So what's the big deal with the output? I'll admit in short documents, it is not easy to see a difference with LaTeX typesetting. However, in longer published works, you begin to see the subtle differences expand dramatically.

Looking closer, you will find LaTeX treats printed output with refined precision. Specifically, the kerning, letter spacing and layout is different from what comes out of a word processor. Consider Figures 1 and 2 from [dartz.free.fr/w/?wakka=latex](http://dartz.free.fr/w/?wakka=latex).

As you can see in Figures 1 and 2, the kerning between characters

### Table

Figure 1. An Example of Microsoft Word Kerning—Incorrect Kerning for the Ta Letter Pair

### Table

Figure 2. The Word Table Processed by LyX/LaTeX—Adjusted Kerning for the Ta Letter Pair

is slightly different. One word does not make a big difference, but a whole page of text does.

### What LyX Does

As you can see from the previous example, LaTeX is ugly to work with in plain-text format. The commands provide fine-looking output, but no one wants to key these in by hand. To fix this problem, several popular LaTeX editing programs are available to do the command formatting for you.

LyX is a GUI document-processing front end for LaTeX. With LyX, you can key in the text and let the program organize how it looks on paper. LyX calls this the What You See Is What You Mean (WYSIWY) way of document processing—meaning you don't need to play with formatting the document. You focus on what you're writing and let the LyX commands do the work of making it look good.

### Where to Get LyX?

LyX is likely in the repository of your Linux distribution. So, all you probably need to do is use your package manager to install the program, and you're ready to begin. If LyX is not in the repository, you can download and install it from [www.lyx.org](http://www.lyx.org).

Besides the LyX package, it's also important to download and install a spell-checking program, such as *ispell* or *aspell*. Again, use your package manager to install these.

### Starting LyX

Prepare yourself—the starting screen of LyX appears stark compared to a typical word processor (Figure 3). Remember, it's not a word processor; it's a text publishing system. LyX won't disappoint you in its capability for delivering good results.

Before going further, set up some defaults applicable to your environment. First, go to Layout→Document (Figure 4). Here you need to define what document you are creating. For this two-part series, we work with the book class in 8.5x11 US Letter. After selecting the book class, select Paper, as shown in Figure 5. Use the drop-down list to find US Letter.

Next, include either *ispell* or *aspell* as the assigned spell checker for LyX. Go to Edit→Preferences, and select Spellchecker under Language Settings. In Figure 6, I have selected *aspell* as the spell checker for LyX.

Finally, review the document converters installed with LyX. Look further down the Preferences screen, and you will see Converters. Select this, and make sure your distribution lists the proper .dvi and .pdf programs in the right location (Figure 7).



# 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](http://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](http://www.pgroup.com) ++ 01 (503) 682-2806

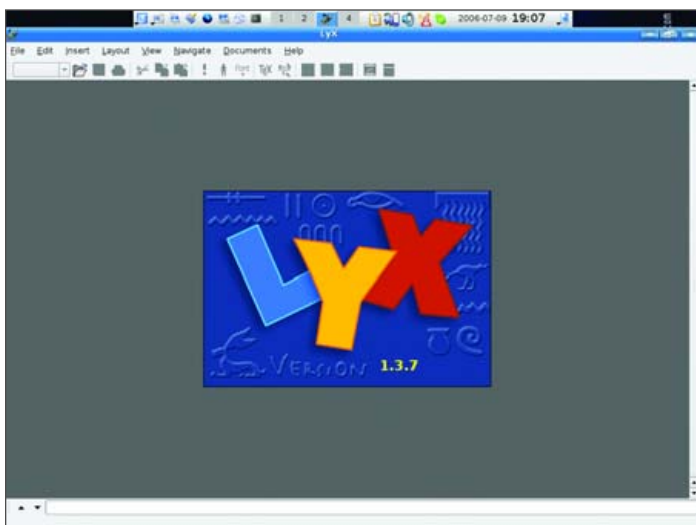


Figure 3. LyX Starting Page

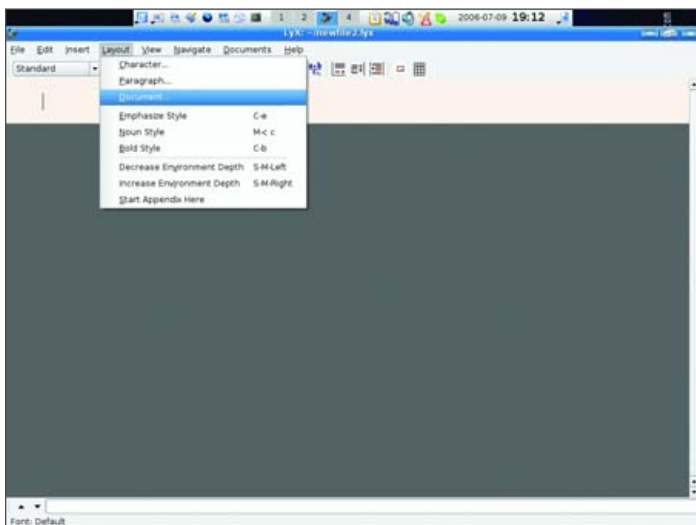


Figure 4. Choose Document to change the characteristics of the whole document.

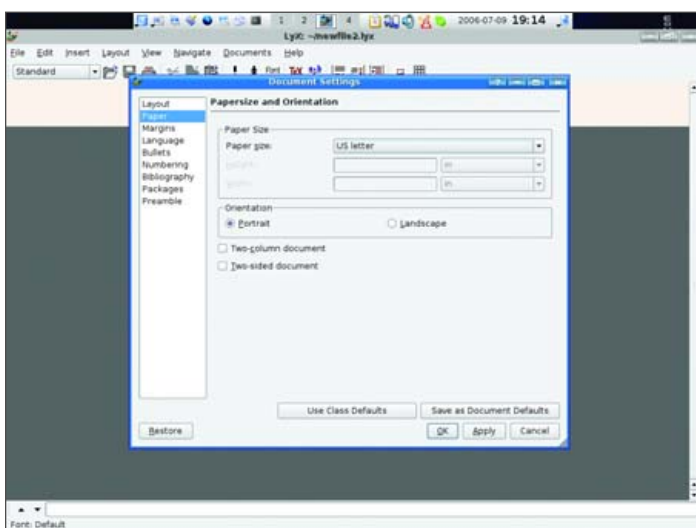


Figure 5. For the right paper size, I chose US Letter.

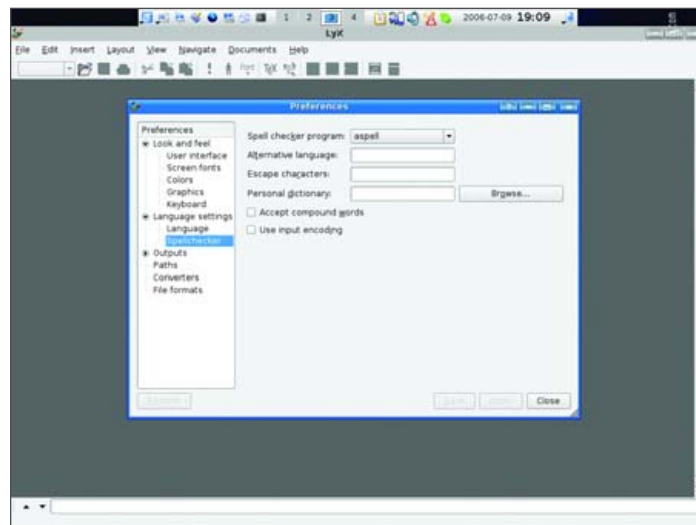


Figure 6. Select aspell or ispell for spell checking.

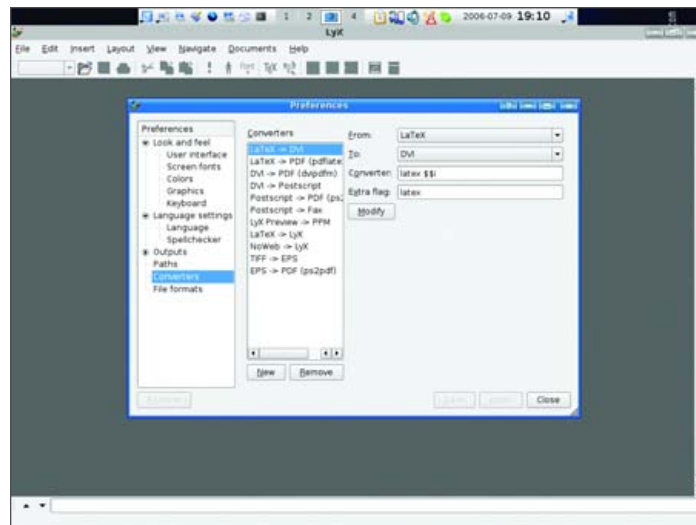


Figure 7. Some distributions (like Kubuntu) have the right converters installed by the package manager.

When finished, you must reconfigure LyX for it to work properly. Go to **Edit**→**Reconfigure**, then restart LyX. Now you're ready to learn this powerful program.

## Writing Your Book

At this point, you're ready to start entering text. Although it's best to get a good understanding of how LyX works first, so you can lay out the final text properly.

Go to **Help**→**Tutorial**, and LyX loads the tutorial into the working screen. Read through the LyX Tutorial, and follow the instructions for creating your first document. The tutorial is easy to understand, and completing the exercises will get you familiar with the program.

I know most of us will prefer the Quick Start Tutorial—so here it is. Go to **File**→**New** and create a new file. Type some sample text on the first line, as shown in Figure 8.

Remember, LyX handles what you want the text to look like on paper based on the assigned document class. So, to create a title page, all you need to do is select **Title** from the drop-down list. LyX marks these words as the book title, and you're done. Now if you press **Ctrl-D**, LyX exports the text into a **.dvi** file and displays your results

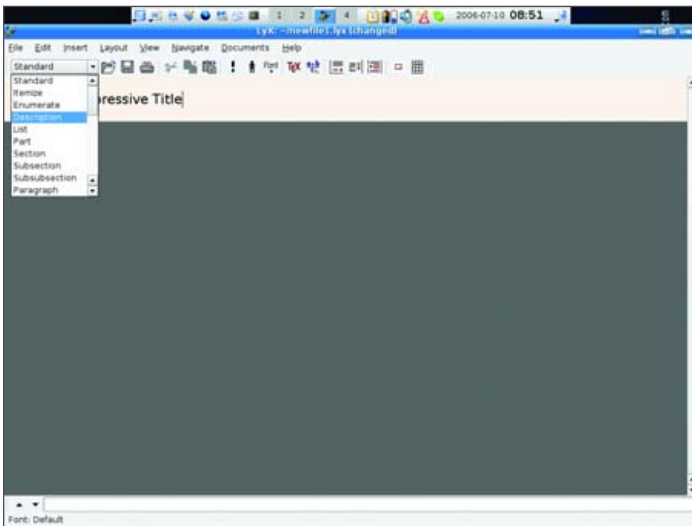


Figure 8. Using the sample text, choose Title to assign the proper format.

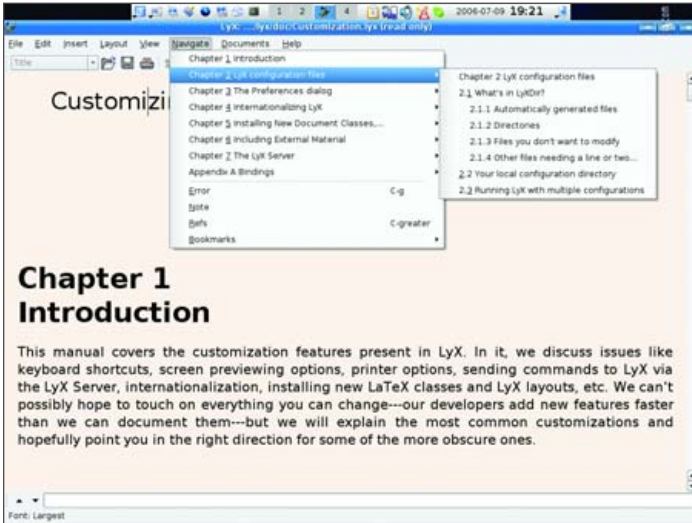


Figure 9. Use the drop-down menus to move through your document quickly.

on-screen. Notice that LyX has centered the text, added the date below the title and turned it into its own separate page—pretty cool.

### LyX Features

Like a word processor, LyX has a few features to help you produce your final work. Academic math people use LyX because it can produce complex formulas fairly easily in printed documents. Doctoral students also use it to conform to standards for their final dissertations.

I find using the mouse awkward when writing. Consequently, I prefer to use command keys and other shortcuts to format text in LyX. LyX comes with a lot of documentation; however, finding answers to your questions can take some looking. A LyX help file titled customization.lyx describes various command keys and bindings to help speed up typing. Print out the file and look through the existing key bindings; they will improve your speed with document processing and keep your focus on what you're writing.

Also, keep in mind that entering a carriage return does not translate into an extra line in LyX. I'll admit, letting the program handle the formatting is unnerving at first, but the results will please you.

LyX automatically creates links to specific parts of the text file. Figure 9 shows how LyX builds a navigation tree based on the text in



# 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 Starts at \$2395**

- 1U Dual Core 5030 CPUs.
- 4GB FBDIMM Memory.
- Supports upto 64GB FBDIMM.
- 120 GB hotswap hard drive.
- 2xIntegrated Dual 10/100 LAN.



**Dual Xeon Server starts at \$4139**

- 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 \$6,699**

- 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](http://www.asacomputers.com)  
 Email: [sales@asacomputers.com](mailto: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.

## TOGETHER, LYX AND LULU MAKE A GREAT PAIR. ALTHOUGH THEY CAN'T FIX YOUR POOR WRITING HABITS, THEY WILL MAKE YOUR FINAL PUBLICATION LOOK PROFESSIONALLY PRINTED AND BOUND.



Figure 10. Here is the same file using Adobe Acrobat.

the document. This is similar to the Outline feature in other word processors and is handy for editing large files.

Like a word processor, LyX handles tables and graphics with ease. To add a table, go to Insert→Tabular Material, and define the table size. Use the Insert drop-down list to place graphics in the document. As an alternative, you can click on the associated icons below the toolbar for graphics, tables and to alter text justification.

Finally, LyX has superior ability to handle cross references, citations and footnotes. While typing, use the Insert drop-down list to add footnotes, citations or cross-reference markers. Each has its own window for the text entered. To keep from viewing them, click the inserted icon, and they disappear off-screen. LyX automatically adjusts the output to keep the footnotes on the proper page.

LyX is powerful, and the documentation is lengthy—too large to cover in this short article. No doubt, using LyX is uncomfortable at first, but the benefits of letting the program sort out the document formatting are profound.

### Lulu—A Self-Publishing Giant

If you haven't heard, Lulu.com ([lulu.com](http://lulu.com)) is a Web site for self-publishing. That's right, you can write your own books, articles, handouts and more. Once it's complete, send your written material to Lulu and select how you want it to be published.

With Lulu, users can choose various publication sizes and bindings. As the author, you can decide whether to glue, staple or stitch the final work. In addition, Lulu even offers hardcover binding.

Further, authors can publish and promote their books directly through Lulu. Lulu has an on-line store where you can browse by subject matter and author. If you want, you even can publish your work as a downloadable book.

As the author, you don't pay anything ahead of time. People who want to buy your work pay a preselected fee, and Lulu takes a cut of the price. Check out Lulu.com for pricing details.

### A Step Above

I think it's likely that most of Lulu's authors use a word processor for their publications. Lulu's help system even provides examples on how to lay out your work from within Microsoft Word and OpenOffice.org. So, using LyX to publish your work will strengthen the professional look of your documents. Thus, even if you write badly, it'll look great.

Lulu.com is a helpful Web site, and authors can upload their work in many different file types. I think the safest way to maintain your work is to use .pdf or .dvi files for upload. This way, you're sure to maintain the nice typesetting look LyX provides.

### The Process

Go to Lulu.com, and sign up for a free account. Then, look through the wide variety of products it offers. For extra fees, you can have Lulu.com help you with the layout of the book and cover design as well.

As a test, I used the LyX tutorial to see how well Lulu works. I exported the file in .pdf format (Figure 10) and used the on-line Lulu instructions to send my file. After Lulu accepts the document, it prompts you to select the binding type, color content and finally the cover design. When you're finished, you can preview the cover of your book and order a copy for final proofreading in the polished format. Each step of the way, Lulu calculates the price of your publication, so you can tailor it for the intended audience.

### Lulu Options

Lulu gives you several choices for publications. You can keep your uploaded documents private for only you to view, or you can release them for public purchase. You decide on the sale price.

As a consultant, I must keep client files and reports so only I can view them. Then, I order just the right amount for the project. After I'm done with the client work, I delete it from Lulu and keep one copy for myself in electronic format.

To get the word out on a self-published book, Lulu offers fee-based services through selected third-party vendors. But wait, this may not be necessary, because in my next article, I'm going to write how to custom create your book cover with Pixel.

### Conclusion

There are many locations on the Internet where you can find tutorials and examples of LyX and LaTeX. Many are difficult to read and understand. My experience is that working with a few documents and following the guidance in the tutorials is enough to get you started with the program.

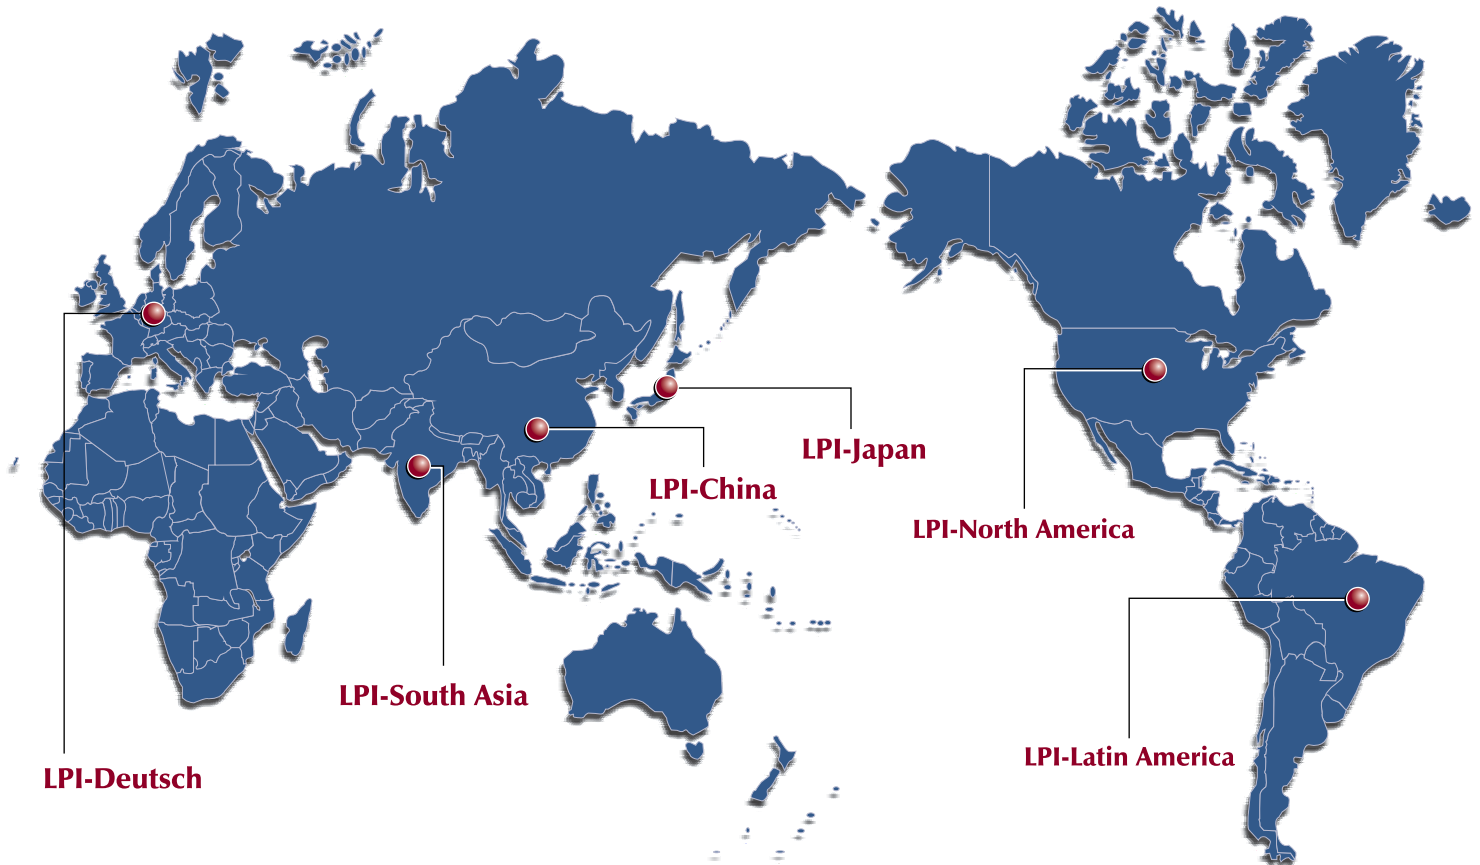
Moreover, since I've been using LyX, I get many comments on how professional the writing looks. As mentioned earlier, the benefits of LaTeX typesetting are really noticeable in larger documents. I think Lulu is a super partner for a good desktop typesetting program. It handles many text formats with ease, and professional binding always looks nice.

So go ahead, write that book you always wanted to write, and make it a best-seller with LyX and Lulu. ■

---

Donald Emmack is Managing Partner of The IntelliGents & Co. He works extensively as a writer and business consultant in North America. You can reach him at [donald@theintelligents.com](mailto:donald@theintelligents.com) or by cruising the 2 meter amateur RF bands in the Midwest.

# 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.**

# Tighter SSH Security with Two-Factor Authentication

How to set up two-factor authentication using a USB pendrive and ssh-agent for root logins.

PAUL SERY

I enthusiastically use two-factor authentication whenever possible, because static passwords aren't the best mechanism around any moat. Traditional passwords are vulnerable to social engineering, key-loggers, yellow post-it notes and—especially as computers become ever faster—to cracking. Tossing them in favor of two-factor authentication is a good idea and helps me sleep better at night.

Unfortunately, network-based, commercial two-factor systems are generally too expensive and complex to use at home or on small networks. But, guess what? You already have the necessary parts on your Linux computer to build a two-factor authentication system. The ubiquitous secure communication tool, OpenSSH, provides all the tools necessary to create a host-based, two-factor authentication system suitable for the home, small office and even larger networks.

This article describes how to combine removable media with OpenSSH public/private keys and the amazing ssh-agent program to achieve two-factor authentication for both regular and privileged users.

## EXAMPLE 1

### Two-Factor User Authentication Using USB Drives

Let's start by creating two-factor authentication for regular (nonroot) users. In this case, we use the well-known SSH public key authentication facility with a small twist. Rather than store the private key in the .ssh subdirectory of your home directory, as is the default, we'll place it on a USB pendrive.

For this example, you'll be logged in as the nonprivileged user bob on a Fedora Core computer, machine1. You'll connect to the remote Linux box machine2 as bob.

Let's start by creating the public/private key pair that we'll use to log in to machine2:

```
ssh-keygen -t rsa -f key-rsa-bob@machine2 -C key-rsa-bob@machine2
```

Enter a passphrase when prompted (the longer and more random the better). By default, the ssh-keygen program creates the key pair in the subdirectory .ssh in your home directory—in this case, /home/bob/.ssh. For this example, I've chosen an arbitrary yet descriptive filename to help identify the intended user and hostname at a glance; this will be important in the succeeding examples, which use multiple keys. (I'm assuming the USB drive is formatted with a Linux filesystem like ext3; vfat works, but you'll need to change the key's file permissions to 400 after every mount.)

Mount your USB pendrive, and you should see it as /media/usbdisk, /media/usbdisk1, /media/disk or /media/disk-1. Move your newly created private key to the appropriate directory and limit access to the owner:

```
mv key-rsa-bob@machine2 /media/usbdisk
chmod 400 /media/usbdisk/key-rsa-bob@machine2
```

Next, copy the public key (key-rsa-bob@machine2.pub) into the /home/bob/.ssh/authorized\_keys file on machine2. Make the

authorized\_keys file readable only by the owner:

```
chmod 400 authorized_keys
```

Now, you can log in to the remote computer, machine2, from machine1, as bob, using the public/private key pairs (the -i option tells the ssh client what key to use):

```
ssh -i /media/usbdisk/key-rsa-bob@machine2 bob@machine2
```

Type in the private key passphrase when prompted, and the OpenSSH server on machine2 logs you in. Unmount and remove the USB device (or removable disc) on machine1, and your private key is protected. You've achieved two-factor authentication: one factor is the key stored on the USB device that you can keep separate from your computer, and the second one is the passphrase you store in your head.

Using SSH public key authentication is a common and familiar process to many. Putting the private key onto removable media is a simple way of physically separating one factor from another.

## EXAMPLE 2

### Two-Factor Root Authentication Using ssh-agent

Example 1 shows how to log in to a remote machine securely using a USB device to separate one authentication factor from another. This works well when logging in as a nonprivileged user but not as root. We have to find a way to log in remotely as the superuser.

One solution would be simply to extend the previous example's method and configure the remote OpenSSH server to allow root logins directly from the network. No passwords or keys will traverse the network, but we would violate the age-old system administration prohibition against directly logging in as root. No shortcuts should be allowed, so we have to figure out how to first log in as a regular user and then as root.

Once again, OpenSSH comes to the rescue. In this case, we continue to use public/private keys but introduce a configuration twist. First, configure the remote SSH service to allow root logins via the internal loopback interface but not the external network. Second, configure the ssh-agent utility to allow the remote machine to authenticate root by querying the keys stored on the local machine.

Here's how the process works:

1. Create a private/public key pair for root on the local machine.
2. Copy the public key into root's authorized\_users file on the remote machine.
3. Run the ssh-add utility locally to cache the private key.
4. ssh to the remote machine and log in as a regular user as described in Example 1; however, this time use the agent-forwarding option.





# Expert Included.

Jon brings years of experience to finding innovative ways to meet his customers' IT challenges. He is a fan of the Rackform iServ R256 with two Dual-Core Intel® Xeon® Processors because he knows that intense computing environments demand powerful, efficient solutions. Servers based on the Dual-Core Intel® Xeon® processor 5100 series offer an unprecedented combination of power and efficiency, including Intel Virtualization Technology, as well as support for future architecture.

When you partner with Silicon Mechanics, you get more than a breakthrough Intel solution - you get an expert like Jon.



5. On the remote machine, ssh to the localhost interface as the root user. The remote OpenSSH daemon queries the local agent, authenticates root, and you can log in as the superuser.

The ssh-agent utility provides just the functionality we're looking for. It allows remote SSH daemons to authenticate users by querying the locally stored cache of decrypted private keys. Keys are never transmitted between machines—the private keys remain stored on removable media on your local workstation.

ssh-agent is powerful, but setting it up can be tricky. First, you need to use the ssh-add utility to decrypt your private key and hand it to ssh-agent. Second, you need to tell ssh-add how to communicate with ssh-agent. ssh-add communicates with ssh-agent via a socket, whose location is stored in the SSH\_AUTH\_SOCK environmental variable. By default, ssh-agent creates sockets with arbitrary names, and setting SSH\_AUTH\_SOCK correctly can take some work.

Fortunately, many Linux distributions, including Fedora Core, automatically set up the necessary ssh-agent/ssh-add connections when you log in graphically (such as on GNOME or KDE). Log in at the console, open a terminal console and type the following:

```
ssh-add -l
```

As long as ssh-add can communicate with ssh-agent, you should see either a list of your public keys or a message like "The agent has no identities".

If, for any reason, ssh-agent isn't running or your SSH\_AUTH\_SOCK variable isn't set, or isn't set correctly, you will get the message "Could not open a connection to your authentication agent". In that case, run the following command:

```
eval `ssh-agent`
```

This starts an ssh-agent instance and automatically sets the environmental variables in your current shell.

Next, create a key pair for root as you did in the first example:

```
ssh-keygen -t rsa -f key-rsa-root@machine2 -C "key-rsa-root@machine2"
```

Move the private key to the removable media and give read access to the owner but nobody else:

```
mv key-rsa-root@machine2 /media/usbdisk  
chmod 400 /media/usbdisk/key-rsa-root@machine2
```

Copy the public key into the /root/.ssh/authorized\_keys file on the remote computer machine2.

Add root's private key on machine2 to ssh-agent by running the

## ssh-add

ssh-add allows you to lock and/or confirm using private keys. Use the -x and -X options to lock and unlock a key. You will create a password to lock the key, and use the password to unlock it. Using the -c option directs ssh-add to prompt you every time ssh-agent is asked to use a key. The prompt is displayed on the machine running ssh-agent and effectively prevents unauthorized users from using your keys.



visit us at [www.siliconmechanics.com](http://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.

following command:

```
ssh-add -t 300 /media/usbdisk/key-rsa-root@machine2
```

Enter the passphrase when prompted, and ssh-agent returns the message "Identity added: key-rsa-root@machine2 (key-rsa-root@machine2)" when it adds the key. (The -t 300 option limits the lifetime of the cache to 300 seconds, or five minutes. Your keys will remain viable forever if you don't specify the lifetime.)

Log in to the remote machine as a regular user:

```
ssh -A -i /media/usbdisk/key-rsa-bob@machine2
```

Enter the passphrase when prompted, and you will log in to machine2. (This command is the same as in Example 1, except we're using the -A option, which turns on agent forwarding.)

Type `ssh-add -l` on machine2, and you should see the root key you just added to ssh-agent. For example:

```
2048 fa:5c:4b:73:88:26:.... /media/usbdisk/key-rsa-root@machine2 (rsa)
```

Next, `su` to root (on machine2), and configure the SSH daemon to allow root logins on the internal loopback interface. Edit the `/etc/ssh/sshd_config` file and add/modify the following options:

```
PermitRootLogin yes
AllowUsers bob@*
AllowUsers root@localhost.*
```

(Some OpenSSH configurations require you to set the numeric loopback address explicitly: `AllowUsers root@127.0.0.1`.)

Save your changes, and restart the SSH daemon:

```
service sshd restart
```

Log out of the root account, and use OpenSSH to log back in as root:

```
ssh root@localhost
```

## Two vs. 2.X Factors

Some people count the locally stored SSH keys and their passphrases as two factors. This view is reasonable, but I feel more comfortable physically separating the key storage device from the computer. Keeping your keys on removable media reduces the opportunity for intruders to capture and crack them.

Now, it's important to realize that keeping your keys on devices like USB pendrives doesn't eliminate the ability of an intruder to spy them. Your keys are vulnerable while mounted, and you should take precautions to harden the workstation from which you connect to other computers. Use good passwords for local (console) logins, keep your workstation patched and so on.

So, you're better off using public key authentication than static passwords, as long as you adequately protect your workstation. How safe you want to be depends on your paranoia.

Now the OpenSSH daemon on machine2 accepts root logins on the loopback interface but not from the external network. It negotiates with ssh-agent on machine1 to authenticate you as the root user. root's private key never left machine1! Using OpenSSH in this way effectively allows you to replace the `su` (switch user) and `sudo` utilities.

But, we're not quite finished. You can increase security further by limiting the `su` command to locally connected devices. Modify `/etc/pam.d/su` as shown below to prevent anyone from using `su` over the network:

```
auth        required    pam_securetty.so
```

The `su` command will work only from the console and virtual terminals.

Unmount and remove your USB device. Individuals actually will have to steal your USB drive at this point to get your keys. Even then, they have to discover your passphrase or expend lots of computing power and time cracking the key.

### EXAMPLE 3 Tightening Up

We need to close a potential vulnerability before using this system in the wild.

Using ssh-agent and agent forwarding allows the remote SSH server to query the private key stored on your local computer. However, if you use this system to log in to multiple computers, an intruder on one machine can potentially hijack those keys to break in to another machine. In that case, this system could be more dangerous than one using static passwords.

To illustrate the problem, let's expand our example network from two to three nodes by adding machine3 to the mix. Create key pairs for both bob and root on machine3, as described in Examples 1 and 2, and add root's private key to ssh-agent on machine1.

Now, `ssh` to machine3 as bob using the agent-forwarding option -A. Run `ssh-add -l`, and you can see the public keys for both machine2 and machine3:

```
2048 fa:5c:4b:73:88:.... /media/usbdisk/key-rsa-root@machine2 (RSA)
2048 26:b6:e3:99:c1:.... /media/usbdisk/key-rsa-root@machine3 (RSA)
```

In this example, ssh-agent on machine1 caches the private keys for machine2 and machine3. This single agent allows us to log in as root on either computer. However, using the single agent also potentially allows an intruder on machine2 to log in as root on machine3 and vice versa. This is not good.

Fortunately, we can fix this problem by using the `ssh-add -c` option; we can add additional security by using individual ssh-agent instances to store one root key for each remote machine. The -c option tells ssh-agent to have the user confirm each use of a cached key. Devoting one ssh-agent instance per host prevents any as yet unknown ssh-agent vulnerability from exposing one machine's key to another.

Using the `ssh-add confirm` option is easy; simply set the -c option whenever adding a key to ssh-agent. Let's give it a try. Start two agents on machine1, specifying predefined sockets:

```
ssh-add -c /media/usbdisk/key-rsa-root@machine2
ssh-add -c /media/usbdisk/key-rsa-root@machine3
```

You'll be asked to confirm use of the key when you `ssh` to machine2 and machine3.

You also can use separate ssh-agents to store each key. Let's give it a try; start two agents on machine1, specifying predefined sockets:

```
ssh-agent -a /tmp/ssh-agent-root@machine2
ssh-agent -a /tmp/ssh-agent-root@machine3
```

Once again, I'm using an arbitrary yet descriptive naming convention. Set the environmental variable, and add the key for machine2:

```
export SSH_AUTH_SOCK=/tmp/ssh-agent-root@machine2
ssh-add -c /media/usbdisk/key-rsa-root@machine2
```

Repeat this process for machine3, making the appropriate substitutions:

```
export SSH_AUTH_SOCK=/tmp/ssh-agent-root@machine3
ssh-add -c /media/usbdisk/key-rsa-root@machine3
```

Now, log in to machine3 (we'll go to machine3 at this point as we just set the SSH\_AUTH\_SOCK variable to point to machine3's agent):

```
ssh -A -i /media/usbdisk/key-rsa-bob@machine2 bob@machine3
```

Run the following command to see what keys you can query on machine1:

```
ssh-add -l
```

You see only the key for root on machine3.

Exit from machine3, change the environmental variable to the machine2 ssh-agent socket, and log in to machine2:

```
export SSH_AUTH_SOCK=/tmp/ssh-agent-root@machine2
ssh -A -i /media/usbdisk/key-rsa-bob@machine2 bob@machine2
```

Check your keys again:

```
ssh-add -l
```

Checking your keys on machine2 and machine3 reveals only the root key for that machine. In the previous example, by using a single ssh-agent, you would have seen the keys for both machine2 and machine3.

Using separate ssh-agent instances for each machine you log in to requires more work.

Resetting the SSH\_AUTH\_SOCK variable every time you want to log in to another machine is impractical. To simplify the process, I've written a simple script `tfssh` (two-factor ssh) to simplify the process. Its syntax is:

```
tfssh [username@]host [keydir]
```

The script [Listing 1 on the LJ FTP site at <ftp.ssc.com/pub/lj/listings/issue152/8957.tgz>] starts ssh-agent when necessary,

## Storing Keys

You can store your keys on any type of removable media. I'm using a USB pendrive in these examples because it's easy to work with and carry around. Feel free to use writable CD-ROMs or DVDs or even floppies if you want.

sets the environmental variable, adds the root keys to ssh-agent and logs in to the remote machine as the user. You also can tell `tfssh` to look in an arbitrary directory ([keydir]) for its keys and also set a key timeout for the key cache.

## Conclusion

Static passwords are quickly becoming more trouble than they're worth. We need to break the static habit and start using two-factor authentication. OpenSSH is a powerful system that provides the tools necessary to make that step. By using public/private keys, agent forwarding and removable media, we can use OpenSSH as a key "safe". This, in turn, allows us to create a simple, inexpensive and effective host-based, two-factor authentication system.

This two-factor system requires a moderate amount of work to configure and use, but it is well worth the extra security. However, using the `tfssh` script makes the process easy to use. Using the script means you get all the benefits of two-factor authentication but almost none of the hassle. ■

---

Paul Sery has been a UNIX and Linux System Administrator for more than 20 years. He's written several Linux books, including *Network Linux Toolkit* and *Knoppix for Dummies*. He's also co-authored several *Red Hat Linux for Dummies* and *Fedora Core for Dummies* books with Jon "maddog" Hall. Paul lives in Albuquerque, New Mexico, and can be reached at [pgsery@swcp.com](mailto:pgsery@swcp.com).

---

**axigen**  
How far can you take your mail server?  
DISCOVER THE NEW  
**AXIGEN MAIL SERVER 2.0**

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

GRAPHICAL ADMINISTRATION INTERFACE (WEB BASED)	ANTIVIRUS & ANTISPAM INTEGRATION (15 APPLICATIONS)	FREE TECHNICAL SUPPORT AROUND THE CLOCK
--	--	---

[www.axigen.com/Lj](http://www.axigen.com/Lj)  
Visit now to start your free trial!

# A Server (Almost) of Your Own

Set up a virtual host for e-mail on your virtual private server. GEORGE BELOTSKY

**Would you like** to have a dedicated server at an ISP, for the price of a mere virtual hosting account? For most Linux users, the answer is certainly, yes. You want root access to your own box and the ability to run the software that you choose—even if the budget calls for virtual hosting.

In this case, the solution is a Virtual Private Server (VPS). VPS accounts effectively partition a physical computer's resources into several virtual machines. You get root access to your VPS and configure it just like you would a dedicated server.

Of course, the flexibility of a VPS comes at the price of increased complexity. You are the system administrator of your VPS, not your ISP. The correct operation of the virtual machine—particularly security—is your responsibility.

The typical VPS account holder, however, needs to support only a small number of users, with a few relatively simple services. This makes the task of administering the system much easier. If you are at least somewhat comfortable with managing a Linux machine from the command line, you should be able to make a successful transition to a VPS account.

In this article, we focus our attention on the most critical aspect of switching to a VPS from virtual hosting—getting your e-mail working. E-mail is one of the most important communication tools today. With the exception of DNS, it is also the most complex service you are likely to encounter. Learning how to get your e-mail working should give you a good overall sense of how to manage your VPS.

With respect to DNS, you may want your VPS provider to handle it for you entirely, at least in the beginning. Ask about the additional fees before you sign up. They should be a few extra dollars per year. Some domain name registrars and third parties also can provide you with DNS service.

## Getting Started

We use the VPS service provided by tummy.com to implement and test our e-mail solution. Its VPS accounts are based on Red Hat's Fedora by default, but you easily can choose Debian instead during the sign-up process. We use the Fedora-based VPS for this article—Fedora Core 3 at the time of writing. Some of the steps shown in the following discussion are specific to Fedora, but most are applicable to any recent Linux distribution. Updates for more recent Linux distributions are available at [www.linuxjournal.com/article/9380](http://www.linuxjournal.com/article/9380).

Here are some names that I use in the examples. Your VPS hostname is myvps, your workstation is ws, your first domain name is first.domain, and your second domain name is second.domain. Your user name on your workstation is usera, and the mail users on the VPS are maila and mailb.

Additional domain names beyond the first one are optional—only remember to delete all references to second.domain when you use any of the code from the article. You also can host more than two domain names—simply configure them in the same way as second.domain is configured in the examples.

Of course, the actual domain names that you use should be valid and registered to you. For example, my first.domain is openlight.com. You also can call your VPS and workstation anything you want. Now, let's begin.

Log in to your new VPS account as root with `ssh root@MY.VPS.IP.ADDRESS`. You would have already chosen your root password when you signed up for the account, and your VPS provider should have given you the IP address of your virtual machine.

One of the first tasks when you set up a new Linux server is to configure the built-in iptables firewall. Your VPS provider may have

set reasonable defaults, but you should always verify this yourself. On the Fedora Linux distribution, run the following command:

```
[root@myvps ~]# system-config-securitylevel-tui
```

You can now move from one control to another with the cursor keys. Use the spacebar to activate buttons and toggle check boxes. Make sure that the Security Level is set to Enabled. Then, activate the Customize button.

On the next screen, you must enable SSH, WWW and Mail. *Do not enable any "Trusted Devices"*.

Next, scroll down to the Other ports text box, and add the entry `https:tcp`, which allows secure https connections. You will need https if you decide to configure Web mail, as described later in this article.

Activate the OK button when you are finished with the customization screen. Also, activate OK on the next screen. Finally, restart iptables to make sure that the changes take effect:

```
[root@myvps ~]# /etc/init.d/iptables restart
```

You must be very careful when you reconfigure your iptables. In addition to the usual danger of creating new vulnerabilities, it is easy to lock yourself out of the remote VPS server. In that situation, you will have to ask your VPS provider for help. Other common ways to render the VPS inaccessible are shutting down networking, the SSH daemon (`sshd`) or halting the virtual machine.

Next, create an ordinary user login that you will use later to read and send e-mail. Set the password for the new account:

```
[root@myvps ~]# useradd maila
[root@myvps ~]# passwd maila
Changing password for user maila.
New UNIX password:
```

Use names such as maila or pseudonyms for logins. This is more secure and guards against inadvertent release of personal information on-line. Verify that you can log in to the new account. You are now ready to configure your mail server.

## WARNING:

There have been many automated attacks against SSH. *At the very least, you must use strong passwords, or your system will be compromised.* The `apg` utility simplifies this task. It generates random, non-dictionary "words" that you can pronounce. There are `apg` packages for most popular Linux distributions.

I strongly recommend that you look in the on-line Resources for this article for more information. SSH security is not specific to the VPS environment, but with a VPS, you do have the flexibility to protect yourself properly.

## Listing 1.

### Example main.cf File for Postfix on a VPS

```
# Note that lines that begin with whitespace
# continue the previous line.
#

# LOCAL PATHNAME INFORMATION
queue_directory = /var/spool/postfix
command_directory = /usr/sbin
daemon_directory = /usr/libexec/postfix

# QUEUE AND PROCESS OWNERSHIP
mail_owner = postfix

# Host name is usually the domain name on a VPS.
myhostname = first.domain
mydomain = first.domain

# Where locally posted mail will come from.
myorigin = $myhostname

# Listen on all interfaces.
inet_interfaces = all

# This server is the final destination for these domains.
mydestination = localhost, localhost.localdomain,
                $myhostname, localhost.$mydomain,
                $mydomain, second.domain

# IMPORTANT -- accept mail for relaying ONLY from
# the local machine.
mynetworks_style = host

# Where your aliases are.
alias_maps = hash:/etc/aliases
alias_database = hash:/etc/aliases

# This user should receive any mail whose recipient
# could not otherwise be matched.
luser_relay = maila@localhost.localdomain

# IMPORTANT -- local recipient checking must be
# turned off for the "luser_relay" directive to
# work.
local_recipient_maps =

# Error code to reject mail with when the local
# recipient is not known.
unknown_local_recipient_reject_code = 550

# Your server's greeting banner. IMPORTANT -- it
# MUST start with your server's hostname, and the
# reverse DNS lookup on the server's IP address MUST
# match this hostname, or your outgoing mail could
# be rejected as SPAM.
smtpd_banner = $myhostname ESMTTP

# See the "main.cf" that came with your Postfix
# distribution for discussion on the rest of the
# directives in this file.
debug_peer_level = 2

debugger_command =
    PATH=/bin:/usr/bin:/usr/local/bin:/usr/X11R6/bin
    xgdb $daemon_directory/$process_name $process_id
    & sleep 5

sendmail_path = /usr/sbin/sendmail.postfix

newaliases_path = /usr/bin/newaliases.postfix

mailq_path = /usr/bin/mailq.postfix

setgid_group = postdrop

html_directory = no

manpage_directory = /usr/share/man

sample_directory = /usr/share/doc/postfix-2.1.5/
                  samples

readme_directory = /usr/share/doc/postfix-2.1.5/
                  README_FILES
```

## Configuring the Mail Server

The mail server, also known as the Mail Transfer Agent (MTA), is a program that delivers and receives e-mail messages. The MTA will receive all the mail that others send you. Likewise, any messages you send to others will leave your VPS through the MTA.

The default MTA on your VPS is Sendmail. This sophisticated, powerful program has advantages for complex e-mail configurations. Unfortunately, it also is difficult to configure and tends to have a lot of security problems.

Therefore, we replace Sendmail with Postfix. Postfix is efficient, very secure and, most important, easy to configure. Before proceeding with the installation, shut down Sendmail, and make sure that it will not start again on reboot. Then, install Postfix:

```
[root@myvps ~]# /etc/init.d/sendmail stop
Shutting down sendmail:          [ OK ]
Shutting down sm-client:        [ OK ]
[root@myvps ~]# chkconfig --del sendmail
[root@myvps ~]# up2date --install postfix
```

Note that using the up2date command to install packages is specific to Red Hat and related distributions. You may be presented with a configuration screen the first time that you run up2date. You can simply press Enter to accept the default values. In addition, up2date is sometimes very slow and can even fail for transient reasons. You can try the command again if it does not work the first time.

The main Postfix configuration file is /etc/postfix/main.cf. Save a copy of this file to read later, because it contains many helpful comments. Then, replace /etc/postfix/main.cf with the code from Listing 1. You should modify your new main.cf to specify the domain names that you will be hosting on your VPS.

Replace all occurrences of first.domain in Listing 1 with your own fully qualified domain name, such as openlight.com. *The reverse DNS lookup of your VPS's IP address must return this domain!* Otherwise, your outbound messages may be rejected as spam.

If you are hosting an additional domain name, substitute it instead of the second.domain entry. Otherwise, delete second.domain before using Listing 1. Also, replace maila in Listing 1 with the user name of your choice.

Now, append an entry to the /etc/aliases file to specify the user

who will receive root's mail. Here is an example:

```
root: maila
```

Next, create accounts for the other e-mail users. Append any aliases for these users to `/etc/aliases`. The following example entry will cause user `mailb` to receive all messages sent to `promo@first.domain`:

```
promo: mailb
```

Note that if you have an additional domain name, messages to `promo@second.domain` will also go to `mailb`. For a small organization, this is probably the right default behavior, because all domain names that you will be hosting are almost certainly related. For example, if you are hosting an additional domain for your product, then tech-support questions about the product should likely go to the same person, regardless of which domain name appears in the e-mail address.

When you are finished, update the alias database file, and start Postfix:

```
[root@myvps ~]# postalias /etc/aliases
[root@myvps ~]# /etc/init.d/postfix start
Starting postfix: [ OK ]
```

Check the log file `/var/log/maillog` for any errors.

You can update the aliases file even while Postfix is running, just

run `postalias /etc/aliases` again when you are finished.

You should now verify that Postfix is doing what you expect. Connect to port 25 on your VPS using Telnet. You can do this by issuing commands interactively to the server, as shown in Listing 2.

Enter the text as shown in Listing 2. Of course, you should type the IP address of your VPS in place of `MY.VPS.IP.ADDRESS`, and your actual domain name instead of `first.domain`. Use Listing 2 as a guide to run the following tests:

- Connect to port 25 of your VPS from an outside machine, such as your workstation. Verify that Postfix accepts messages for each e-mail address you intend to use. Then, make sure the right users are receiving the messages. See the following discussion for details.
- Connect again from the outside, and check that Postfix will refuse to relay mail to other systems. Use an e-mail account that you have on some other system as the destination, just in case. *It is very important that your MTA refuse any relay requests from external machines.* Otherwise, spammers quickly will discover that they can route their junk e-mails through your system.
- Using Telnet from a shell prompt on your VPS itself, check that the MTA will relay mail to remote servers. Use your own e-mail account on some other system as the destination. Note that the remote MTA may refuse to accept the message, because your system is not live yet, so reverse DNS lookups will not yield the right result.

You can verify that a user on the VPS has received mail with the `mail` command. Here is an example that checks the mail for `maila`:

```
[root@myvps ~]# mail -u maila
```

The `mail` command is a simple mail reader. Type `h` to view the received messages, then type the number of the message to view it. Press the spacebar to scroll through the message. You also can scroll through the message with the Enter key, but it will start viewing the next message after it gets to the end of the current one. You can type `q` to stop viewing a message. When you are not viewing a message, typing `q` will exit mail. The `?` key brings up a help screen.

When everything is working as it should, tell your initialization scripts to launch Postfix automatically on system reboot:

```
[root@myvps ~]# chkconfig --add postfix
```

### Preparing to Read Your Mail

In this article, we discuss two methods for reading your mail. One is to download the mail to your workstation. The other is to leave it on the VPS and use a Web-based solution to view the messages through your browser. You can use both methods together.

The first approach requires the POP3 protocol, and the second needs IMAP. On Fedora, the simplest way to get both is to install `dovecot`:

```
[root@myvps ~]# up2date --install dovecot
```

When the installation finishes, edit `/etc/dovecot.conf`. Find the `protocols` directive and replace it with the following. Do not delete the original line, but comment it out for future reference:

```
#protocols = imap imaps pop3 pop3s
protocols = pop3 imap
```

As a security precaution, configure both POP3 and IMAP to accept

**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](http://www.offassist.com/linuxjournal)

[www.offassist.com](http://www.offassist.com)  
We Speak Geek! Open Source — Open Mind

## Listing 2.

### Verifying that Postfix Is Working Properly

```
[usera@ws]$ telnet MY.VPS.IP.ADDRESS 25
Trying MY.VPS.IP.ADDRESS
Connected to MY.VPS.IP.ADDRESS.
Escape character is '^]'.
220 first.domain ESMTP
HELO example.com
250 first.domain
MAIL FROM: test@example.com
250 Ok
RCPT TO: promo@first.domain
250 Ok
DATA
354 End data with <CR><LF>.<CR><LF>
This is a test

250 Ok: queued as MESSAGEID
QUIT
221 Bye
Connection closed by foreign host.
```

requests only from the VPS itself. Once again, do not delete the original code, but leave it commented out for future reference:

```
#imap_listen = [::] imap_listen = [127.0.0.1] #pop3_listen = [::]
pop3_listen = [127.0.0.1]
```

Start dovecot, and add it to your system's initialization scripts:

```
[root@myvps ~]# /etc/init.d/dovecot start
Starting Dovecot Imap: [ OK ]
[root@myvps ~]# chkconfig --level 345 dovecot on
```

### How to Read and Send Mail from Your Workstation

We will be using SSH tunneling to read and send mail securely from your workstation. With SSH tunneling, you can temporarily map ports on the VPS to available ports on the workstation. All communication on the mapped ports takes place over an encrypted tunnel.

Give the following command on your workstation. Use your VPS's IP address if you did not add an entry for myvps in the /etc/hosts file on your workstation:

```
[usera@ws ~]$ ssh -Nf maila@myvps \ -L 2525:localhost:25 -L
2110:localhost:110
```

The user maila must have shell access to the VPS. You will be prompted for maila's password.

This tunnel maps ports 25 and 110 on the VPS to ports 2525 and 2110 on the workstation, respectively. If you are already downloading your inbound mail using POP3 and sending your outbound mail through an ISP's mail server, you will require very few changes to your mail client's configuration.

Simply set your POP3 server as localhost with port 2110, and your outbound mail server to localhost with port 2525. You even can leave your outbound mail settings unchanged, unless you plan to cancel the account at the ISP whose mail server you are currently using. Here are specific instructions for two popular e-mail clients.

If you use Mozilla Thunderbird, select Account Settings... from the Edit menu. Add a new account by clicking the Add Account... button in the dialog box, and follow the prompts in the Account Wizard. After you create the new account, click on its Server Settings list item in the



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  
Upto 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](mailto: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](mailto: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](mailto:sales@genstor.com)

Fax: 1-408-383-0121

Prices, Products and Availability subject to change without notice.

left pane to configure the POP3 server and port. Figure 1 shows the screenshot. I have highlighted the most important parameters in red.

You also can configure Thunderbird's outgoing mail server from the same Account Settings dialog box. Click on Outgoing Server (SMTP) in the left pane of the dialog. Figure 2 shows the resulting screenshot. Remember to uncheck the Use name and password check box.

Another popular mail client is Mutt. A typical Mutt-based configuration uses fetchmail to download the mail, procmail to sort it into mailboxes and ssmtp to deliver the outbound mail. See Listing 3 for an example .fetchmailrc file and Listing 4 for an example ssmtp.conf file. Both use the SSH tunnel that we created earlier. Do not forget to change the code in Listing 3 to reflect your correct user names and passwords.

Finally, note that you need to set up the SSH tunnel again every time you reboot your workstation. There are many ways to automate the process, but it is beyond the scope of this article to discuss them.

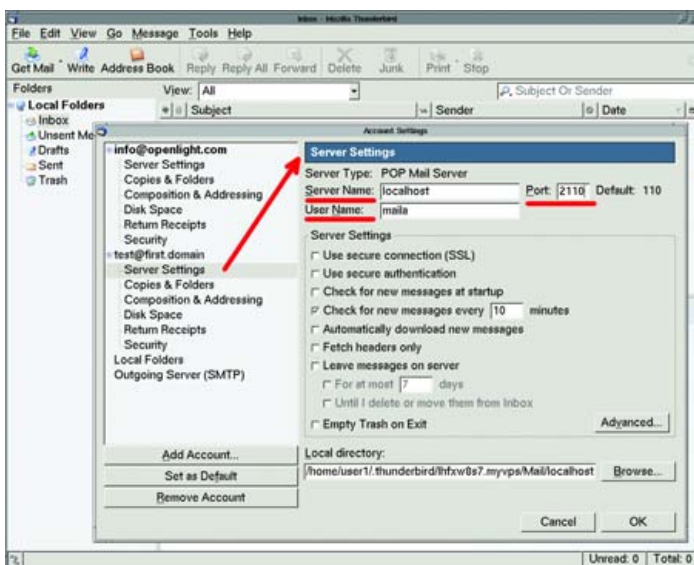


Figure 1. Setting Up Your Mail Account in Mozilla Thunderbird

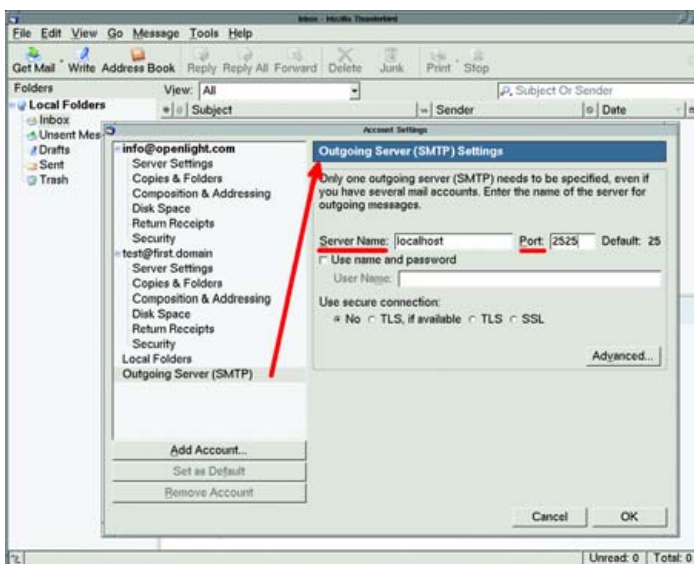


Figure 2. Set up your outgoing mail server to localhost at port 2525.

## How to Read and Send Mail over the Web

The Fedora Linux distribution provides a Web-based e-mail interface that requires very little work to configure. It is based on SquirrelMail and Apache. Web mail is an easy way to support Windows clients. It also does not require shell access on the VPS.

First, install SquirrelMail:

```
[root@myvps ~]# up2date --install squirrelmail
```

This process also installs several other packages that SquirrelMail requires. Next, enable secure https access by installing mod\_ssl:

```
[root@myvps ~]# up2date --install mod_ssl
```

You must disable unsecure http access to SquirrelMail. Edit the file /etc/httpd/conf.d/squirrelmail.conf, and append the following lines:

```
<LocationMatch "/webmail">
  SSLRequireSSL
</LocationMatch>
```

Now, start the Apache Web server:

```
[root@myvps ~]# /etc/init.d/httpd start
```

### Listing 3. The .fetchmailrc Configuration File

```
set postmaster "usera"
set no bouncemail
set no spambounce

poll localhost with protocol POP3 and port 2110
and options no dns:
  user "maila" there is usera here and wants
  mda "/usr/bin/procmail -d %T" options fetchall
  password "MAILA'S VPS PASSWORD"
```

### Listing 4. The ssmtp.conf Configuration File

```
# The person who gets all mail for userids < 1000
# Make this empty to disable rewriting.
root=postmaster

# The place where the mail goes. The actual machine
# name is required; no MX records are consulted.
mailhub=localhost:2525

# The full hostname
hostname=localhost

# The "From" line sender address will override any
# settings here.
FromLineOverride=YES
```



Connect to `https://MY.VPS.IP.ADDRESS/webmail`. Your browser will warn you about the SSL certificate—just accept it permanently, and you will not be warned again. The only way to avoid this error altogether is to use a certificate signed by a recognized Certificate Authority (CA). The CA will need to verify your identity and also will charge an annual fee for signing the certificate.

After accepting the certificate, you should be able to log in as any of the mail users that you have created earlier. If a particular mail user—for example `mailb`—does not need shell access, disable it with the following command:

```
[root@myvps ~]# usermod -s /sbin/nologin mailb
```

Do not forget to add the Apache Web server to your startup environment:

```
[root@myvps ~]# chkconfig --level 345 httpd on
```

Your Web mail users should click on the Options link in the SquirrelMail interface and configure their account information. Otherwise, SquirrelMail will format their messages with something like `mailb@localhost.localdomain` in the From field. This certainly will confuse anyone who receives such a message.

## Conclusion

This article has covered one of the most difficult aspects of switching to a VPS account—setting up your e-mail. As you have seen, e-mail service is provided by a collection of several different programs working together. There are many other ways to configure this service. Unfortunately, it would require a lengthy book to describe and compare them all. This article tries to provide a simple solution with good security that a new VPS user can implement quickly.

Welcome to the world of VPS hosting—the server that is (almost) your own.

## Acknowledgements

The author wishes to acknowledge Sean Reifschneider and Evelyn Mitchell of `tummy.com, LTD.`, for generously providing a VPS account used to test the examples presented here, as well as their valuable comments on this article. ■

Resources for this article: [www.linuxjournal.com/article/9380](http://www.linuxjournal.com/article/9380).

George Belotsky is a software architect who has done extensive work on high-performance Internet servers, as well as hard real-time and embedded systems. His technology interests include C++, Python and Linux. He is also the author of the Flightdeck-UI Open Source Project, which uses the ideas from aircraft instrumentation to implement computer user interfaces. You can reach George at [questions@openlight.com](mailto:questions@openlight.com).

# Don't Gamble with Other Service Providers...Host with Hurricane Electric and Get **Fast, Reliable** Service!

## 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](mailto:sales@he.net) or call 510.580.4190



\*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

# Examining Load Average

Understanding workload averages as opposed to CPU usage. RAY WALKER

**Many Linux administrators** and support technicians regularly use the top utility for real-time monitoring of their system state. In some shops, it is very typical to check top first when there is any sign of trouble. In that case, top becomes the de facto critical measurement of the machine's health. If top looks good, there must not be any system problems. top is rich with information—memory usage, kernel states, process priorities, process owner and so forth all can be obtained from top. But, what is the purpose of those three curious load averages, and what exactly are they trying to tell me? To answer those questions, an intuitive as well as a detailed understanding of how the values are formed are necessary. Let's start with intuition.

## The Intuitive Interpretation

The three load-average values in the first line of top output are the 1-minute, 5-minute and 15-minute average. (These values also are displayed by other commands, such as uptime, not only top.) That means, reading from left to right, one can examine the aging trend and/or duration of the particular system state. The state in question is CPU load—not to be confused with CPU percentage. In fact, it is precisely the CPU load that is measured, because load averages do not include any processes or threads waiting on I/O, networking, databases or anything else not demanding the CPU. It narrowly focuses on what is actively demanding CPU time. This differs greatly from the CPU percentage. The CPU percentage is the amount of a time interval (that is, the sampling interval) that the system's processes were found to be active on the CPU. If top reports that your program is taking 45% CPU, 45% of the samples taken by top found your process active on the CPU. The rest of the time your application was in a wait. (It is important to remember that a CPU is a discrete state machine. It really can be at only 100%, executing an instruction, or at 0%, waiting for something to do. There is no such thing as using 45% of a CPU. The CPU percentage is a function of time.) However, it is likely that your application's rest periods include waiting to be dispatched on a CPU and not on external devices. That part of the wait percentage is then very relevant to understanding your overall CPU usage pattern.

The load averages differ from CPU percentage in two significant ways: 1) load averages measure the trend in CPU utilization not only an instantaneous snapshot, as does percentage, and 2) load averages include all demand for the CPU not only how much was active at the time of measurement.

Authors tend to overuse analogies and sometimes run the risk of either insulting the reader's intelligence or oversimplifying the topic to the point of losing important details. However, freeway traffic patterns are a perfect analogy for this topic, because this model encapsulates the essence of resource contention and is also the chosen metaphor by many authors of queuing theory books. Not surprisingly, CPU contention is a queuing theory problem, and the concepts of arrival rates, Poisson theory and service rates all apply. A four-processor machine can be visualized as a four-lane freeway. Each lane provides the path on which instructions can execute. A vehicle can represent those instructions. Additionally, there are vehicles on the entrance lanes ready to travel down the freeway, and the four lanes either are ready to accommodate that demand or they're not. If all freeway lanes are

jammed, the cars entering have to wait for an opening. If we now apply the CPU percentage and CPU load-average measurements to this situation, percentage examines the relative amount of time each vehicle was found occupying a freeway lane, which inherently ignores the pent-up demand for the freeway—that is, the cars lined up on the entrances. So, for example, vehicle license XYZ 123 was found on the freeway 30% of the sampling time. Vehicle license ABC 987 was found on the freeway 14% of the time. That gives a picture of how each vehicle is utilizing the freeway, but it does not indicate demand for the freeway.

Moreover, the percentage of time these vehicles are found on the freeway tells us nothing about the overall traffic pattern except, perhaps, that they are taking longer to get to their destination than they would like. Thus, we probably would suspect some sort of a jam, but the CPU percentage would not tell us for sure. The load averages, on the other hand, would.

This brings us to the point. It is the overall traffic pattern of the freeway itself that gives us the best picture of the traffic situation, not merely how often cars are found occupying lanes. The load average gives us that view because it includes the cars that are queuing up to get on the freeway. It could be the case that it is a nonrush-hour time of day, and there is little demand for the freeway, but there just happens to be a lot of cars on the road. The CPU percentage shows us how much the cars are using the freeway, but the load averages show us the whole picture, including pent-up demand. Even more interesting, the more recent that pent-up demand is, the more the load-average value reflects it.

Taking the discussion back to the machinery at hand, the load averages tell us by increasing duration whether our physical CPUs are over- or under-utilized. The point of perfect utilization, meaning that the CPUs are always busy and, yet, no process ever waits for one, is the average matching the number of CPUs. If there are four CPUs on a machine and the reported one-minute load average is 4.00, the machine has been utilizing its processors perfectly for the last 60 seconds. This understanding can be extrapolated to the 5- and 15-minute averages.

In general, the intuitive idea of load averages is the higher they rise above the number of processors, the more demand there is for the CPUs, and the lower they fall below the number of processors, the more untapped CPU capacity there is. But all is not as it appears.

## The Wizard behind the Curtain

The load-average calculation is best thought of as a moving average of processes in Linux's run queue marked running or uninterruptible. The words "thought of" were chosen for a reason: that is how the measurements are meant to be interpreted, but not exactly what happens behind the curtain. It is at this juncture in our journey when the reality of it all, like quantum mechanics, seems not to fit the intuitive way as it presents itself.

The load averages that the top and uptime commands display are obtained directly from /proc. If you are running Linux kernel 2.4 or later, you can read those values yourself with the command `cat /proc/loadavg`. However, it is the Linux kernel that produces those values in /proc. Specifically, timer.c and sched.h work together to do the computation. To understand what timer.c does for a

living, the concept of time slicing and the jiffy counter help round out the picture.

In the Linux kernel, each dispatchable process is given a fixed amount of time on the CPU per dispatch. By default, this amount is 10 milliseconds, or 1/100th of a second. For that short time span, the process is assigned a physical CPU on which to run its instructions and allowed to take over that processor. More often than not, the process will give up control before the 10ms are up through socket calls, I/O calls or calls back to the kernel. (On an Intel 2.6GHz processor, 10ms is enough time for approximately 50-million instructions to occur. That's more than enough processing time for most application cycles.) If the process uses its fully allotted CPU time of 10ms, an interrupt is raised by the hardware, and the kernel regains control from the process. The kernel then promptly penalizes the process for being such a hog. As you can see, that time slicing is an important design concept for making your system seem to run smoothly on the outside. It also is the vehicle that produces the load-average values.

The 10ms time slice is an important enough concept to warrant a name for itself: quantum value. There is not necessarily anything inherently special about 10ms, but there is about the quantum value in general, because whatever value it is set to (it is configurable, but 10ms is the default), it controls how often at a minimum the kernel takes control of the system back from the applications. One of the many chores the kernel performs when it takes back control is to increment its jiffies counter. The jiffies counter measures the number of quantum ticks that have occurred since the system was booted. When the quantum timer pops, timer.c is entered at a function in the kernel called timer.c:do\_timer(). Here, all interrupts are disabled so the code is not working with moving targets. The jiffies counter is incremented by 1, and the load-average calculation is checked to see if it should be computed. In actuality, the load-average computation is not truly calculated on each quantum tick, but driven by a variable value that is based on the HZ frequency setting and tested on each quantum tick. (HZ is not to be confused with the processor's MHz rating. This variable sets the pulse rate of particular Linux kernel activity and 1HZ equals one quantum or 10ms by default.) Although the HZ value can be configured in some versions of the kernel, it is normally set to 100. The calculation code uses the HZ value to determine the calculation frequency. Specifically, the timer.c:calc\_load() function will run the averaging algorithm every 5 \* HZ, or roughly every five seconds. Following is that function in its entirety:

```
unsigned long avenrun[3];

static inline void calc_load(unsigned long ticks)
{
    unsigned long active_tasks; /* fixed-point */
    static int count = LOAD_FREQ;

    count -= ticks;
    if (count < 0) {
        count += LOAD_FREQ;
        active_tasks = count_active_tasks();
        CALC_LOAD(avenrun[0], EXP_1, active_tasks);
        CALC_LOAD(avenrun[1], EXP_5, active_tasks);
        CALC_LOAD(avenrun[2], EXP_15, active_tasks);
    }
}
```

The avenrun array contains the three averages we have been discussing. The calc\_load() function is called by update\_times(), also found in timer.c, and is the code responsible for supplying the calc\_load()

## Advertiser Index

For advertising information, please contact our sales department at 206-782-7733 ext. 2 or ads@ssc.com. [www.linuxjournal.com/advertising](http://www.linuxjournal.com/advertising)

Advertiser	Page #	Advertiser	Page #
ABERDEEN, LLC <a href="http://www.aberdeeninc.com">www.aberdeeninc.com</a>	47	MBX <a href="http://www.mbx.com">www.mbx.com</a>	21
AMAX <a href="http://www.amax.com">www.amax.com</a>	43	MICROWAY, INC. <a href="http://www.microway.com">www.microway.com</a>	C4
AML <a href="http://www.amltd.com">www.amltd.com</a>	48	MIKRO TIK <a href="http://www.routerboard.com">www.routerboard.com</a>	7
APPRO HPC SOLUTIONS <a href="http://appro.com">appro.com</a>	C2	MONARCH COMPUTERS <a href="http://www.monarchcomputer.com">www.monarchcomputer.com</a>	10, 11
ASA COMPUTERS <a href="http://www.asacomputers.com">www.asacomputers.com</a>	27, 67	OFFASSIST <a href="http://www.offassist.com">www.offassist.com</a>	76
AVOCENT CORPORATION <a href="http://www.avocent.com/special">www.avocent.com/special</a>	1	OPEN SOURCE SYSTEMS, INC. <a href="http://www.opensourcesystems.com">www.opensourcesystems.com</a>	13
CARI.NET <a href="http://www.cari.net">www.cari.net</a>	85	OPERA SOFTWARE <a href="http://www.opera.com">www.opera.com</a>	19
CIARA TECHNOLOGY <a href="http://www.ciara-tech.com">www.ciara-tech.com</a>	8, 9	PENGUIN COMPUTING <a href="http://www.penguincomputing.com">www.penguincomputing.com</a>	31
CORAID, INC. <a href="http://www.coraid.com">www.coraid.com</a>	53	POGO LINUX <a href="http://www.pogolinux.com">www.pogolinux.com</a>	5
COYOTE POINT <a href="http://www.coyotepoint.com">www.coyotepoint.com</a>	3	POLYWELL COMPUTERS, INC. <a href="http://www.polywell.com">www.polywell.com</a>	83
EMAC, INC. <a href="http://www.emacinc.com">www.emacinc.com</a>	40	THE PORTLAND GROUP <a href="http://www.pggroup.com">www.pggroup.com</a>	65
EMPERORLINUX <a href="http://www.emperorlinux.com">www.emperorlinux.com</a>	15	RACKSPACE MANAGED HOSTING <a href="http://www.rackspace.com">www.rackspace.com</a>	C3
ETNUS <a href="http://www.etnus.com">www.etnus.com</a>	63	R CUBED TECHNOLOGIES <a href="http://www.rcubedtech.com">www.rcubedtech.com</a>	49
FAIRCOM CORPORATION <a href="http://www.faircom.com">www.faircom.com</a>	29	SERVERS DIRECT <a href="http://www.serversdirect.com">www.serversdirect.com</a>	39
GECAD TECHNOLOGIES/AXIGEN <a href="http://www.axigen.com">www.axigen.com</a>	73	SILICON MECHANICS <a href="http://www.siliconmechanics.com">www.siliconmechanics.com</a>	51, 71
GENSTOR SYSTEMS, INC. <a href="http://www.genstor.com">www.genstor.com</a>	77	SPIDERTOOLS.COM <a href="http://www.spidertools.com">www.spidertools.com</a>	35
HURRICANE ELECTRIC <a href="http://www.he.net">www.he.net</a>	79	SUPERMICRO <a href="http://www.supermicro.com">www.supermicro.com</a>	41
INTEL <a href="http://www.intel.com/software/products">www.intel.com/software/products</a>	25	SWELL SOFTWARE, INC. <a href="http://www.swellsoftware.com">www.swellsoftware.com</a>	82
IRON SYSTEMS <a href="http://www.ironsystems.com">www.ironsystems.com</a>	91	TEAMHPC <a href="http://www.teamhpc.com">www.teamhpc.com</a>	37
LANTRONIX <a href="http://www.lantronix.com">www.lantronix.com</a>	45	TECHNOLOGIC SYSTEMS <a href="http://www.embeddeddx86.com">www.embeddeddx86.com</a>	22
LAYER 42 NETWORKS <a href="http://www.layer42.net">www.layer42.net</a>	87	THINKMATE <a href="http://www.thinkmate.com">www.thinkmate.com</a>	17
LINUX JOURNAL <a href="http://www.linuxjournal.com">www.linuxjournal.com</a>	6, 95	@XI COMPUTER CORPORATION <a href="http://www.xicomputer.com">www.xicomputer.com</a>	23
LPI <a href="http://www.lpi.org">www.lpi.org</a>	69	ZT GROUP INTERNATIONAL <a href="http://www.ztgroup.com">www.ztgroup.com</a>	33

function with the ticks parameter. Unfortunately, this function does not reveal its most interesting aspect: the computation itself. However, that can be located easily in sched.h, a header used by much of the kernel code. In there, the CALC\_LOAD macro and its associated values are available:

```
extern unsigned long avenrun[]; /* Load averages */

#define FSHIFT 11 /* nr of bits of precision */
#define FIXED_1 (1<<FSHIFT) /* 1.0 as fixed-point */
#define LOAD_FREQ (5*HZ) /* 5 sec intervals */
#define EXP_1 1884 /* 1/exp(5sec/1min) as fixed-point */
#define EXP_5 2014 /* 1/exp(5sec/5min) */
#define EXP_15 2037 /* 1/exp(5sec/15min) */

#define CALC_LOAD(load,exp,n) \
    load *= exp; \
    load += n*(FIXED_1-exp); \
    load >>= FSHIFT;
```

Here is where the tires meet the pavement. It should now be evident that reality does not appear to match the illusion. At least, this is certainly not the type of averaging most of us are taught in grade school. But it is an average nonetheless. Technically, it is an exponential

decay function and is the moving average of choice for most UNIX systems as well as Linux. Let's examine its details.

The macro takes in three parameters: the load-average bucket (one of the three elements in avenrun[]), a constant exponent and the number of running/uninterruptible processes currently on the run queue. The possible exponent constants are listed above: EXP\_1 for the 1-minute average, EXP\_5 for the 5-minute average and EXP\_15 for the 15-minute average. The important point to notice is that the value decreases with age. The constants are magic numbers that are calculated by the mathematical function shown below:

$$y = \frac{2^{11}}{2^{((5 \log_2(e))/60x)}}$$

When  $x=1$ , then  $y=1884$ ; when  $x=5$ , then  $y=2014$ ; and when  $x=15$ , then  $y=2037$ . The purpose of the magical numbers is that it allows the CALC\_LOAD macro to use precision fixed-point representation of fractions. The magic numbers are then nothing more than multipliers used against the running load average to make it a moving average. (The mathematics of fixed-point representation are beyond the scope of this article, so I will not attempt an explanation.) The purpose of the exponential decay function is that it not only smooths the dips and spikes by maintaining a useful trend line, but it accurately decreases the quality of what it measures

as activity ages. As time moves forward, successive CPU events increase their significance on the load average. This is what we want, because more recent CPU activity probably has more of an impact on the current state than ancient events. In the end, the load averages give a smooth trend from 15 minutes through the current minute and give us a window into not only the CPU usage but also the average demand for the CPUs. As the load average goes above the number of physical CPUs, the more the CPU is being used and the more demand there is for it. And, as it recedes, the less of a demand there is. With this understanding, the load average can be used with the CPU percentage to obtain a more accurate view of CPU activity.

It is my hope that this serves not only as a practical interpretation of Linux's load averages but also illuminates some of the dark mathematical shadows behind them. For more information, a study of the exponential decay function and its applications would shed more light on the subject. But for the more practical-minded, plotting the load average vs. a controlled number of processes (that is, modeling the effects of the CALC\_LOAD algorithm in a controlled loop) would give you a feel for the actual relationship and how the decaying filter applies. ■

Ray Walker is a consultant specializing in UNIX kernel-level code. He has been a software developer for more than 25 years, working with Linux since 1995. He can be contacted at ray.rwalk2730@gmail.com.

# PEG<sup>®</sup>

the leading GUI for  
Embedded Systems

Prototype Today!!  
FREE EVALUATION PACKAGE

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





**SWELL**  
software

GRAPHICS SOFTWARE FOR EMBEDDED SYSTEMS

WWW.SWELLSOFTWARE.COM

810-982-5955

# 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® Optreron™ Dual-Core Processors 2218+
- with Hyper Transport Technology
- Upto 64GB 667/533MHz ECC DDR2 (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
- SATA RAID-5 or SCSI Controller on-board
- 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



1114SAS-2500A, 2218+, 32GB, 2TB **\$8,599**

1114SAS-2500A, 2212+, 16GB, 1TB **\$4,399**

1112ES-2500A, 2210E, 2GB, 500GB **\$1,999**



4024AIS 4U 12TB SAS/SATA **\$10,500**

4024AIS 4U 6TB StorageServer **\$6,399**

#### 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® Optreron™ 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
- 2x320G 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 **\$20,999**

32G RAM, 2x265+, QuadroFX3450 **\$8,499**

16G RAM, 2x246+, QuadroFX1500 **\$4,650**

DDR2 Socket-F is now available

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

#### 1U Power Saving ISP Server

- AMD Sempron™ or Optreron™ 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



#### 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

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

- 8 or 4 AMD® Optreron™ Dual-Core Processors with Hyper Transport Technology
- Upto 128GB DDR Memory for 16-way (32 sockets)
- Upto 64GB DDR2 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 8212, 8GB, 900GB, 8450A **\$7,499**

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](http://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  
Optreron, 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.

# Thinking Thin

Connecting thin clients to Linux. LYLE FROST

**Once upon a time**, there was the mainframe. All application processing was centralized to this enormous beast, and desktop equipment did nothing but display its output. Then the personal computer arrived, ending the tyranny of the mainframe. Individual users suddenly were empowered to install their own applications. Software development and innovation boomed. The personal computers were networked. Thus, the mainframe was slain.

But all did not live happily ever after. The cost of maintaining a workstation on every desktop outgrew the purchase cost long ago. The fact that the dominant operating system is like a Petri dish for viruses and spyware has exacerbated the situation to a point that should be considered intolerable. It has to be faced that, in most situations, it is not desirable to allow the user to install software. The only sane management decision is to draw a clear line between users and administrators.

This can be accomplished in large part by using a secure system like Linux on the desktop. Viruses and spyware disappear, and maintenance costs can plummet. But, there is still a full system on every desktop that must be maintained. Hard drives fail. Fans fail. Major OS updates are not automatic. Desk space is consumed.

One solution is a step forward that feels like turning back the clock. The thin client is the modern equivalent of the text terminal. It provides a low-profile, low-maintenance appliance for the desktop. Application processing is off-loaded to a centralized system called a terminal server. Linux has emerged as the OS of choice on the thin client, even when the terminal server runs MS Windows. But let's not go halfway. Let's explore in detail how to deploy a Linux thin client with a Linux terminal server.

## The Thin Client

What makes a client thin? Most important, thin clients have minimal local software that can be stored on a Flash memory module that is read-only for the local user. This is usually a standard CompactFlash card or a Disk On Module (DOM), which is Flash memory with an IDE interface. A small portion of Flash is made writable for saving configuration information, but in a properly configured system, the user will not be able to modify this. Once configured, it is very nearly an appliance as far as the user is concerned.

Because most of the processing is performed by the terminal server, a slower CPU can be used; 533MHz is typical. This diminishes the cooling requirements greatly, which means fewer or no fans. The silence is golden.

Because there are no internal drives or expansion cards, motherboard components are reduced, allowing very small form factors. The small form factor, reduced cooling requirements and lack of drives mean a very small enclosure. The model I typically use measures 9.5" tall and 1.75" wide, and has a maximum power consumption of 30W. The smaller power supply also means a smaller UPS. Compare a 700 VA workstation UPS costing \$120 US and weighing 17 pounds to a 350 VA thin-client UPS costing \$40 US and weighing 11 pounds.

Thin clients have two distinct modes of operation: client and stand-alone. In standalone mode, the thin client isn't really a client. All necessary applications are loaded in Flash and executed locally, which can drive the purchase cost up by increasing the Flash requirements. The most common application of this is a Web appliance. Any decent thin client will have the ability to boot directly into a Web browser and even prevent



Figure 1. Igel 364 LX



Figure 2. Igel 364 LX Internals

the user from exiting the browser or modifying its configuration.

Here is a big caveat to thin clients: vendor dependence. You can't simply download the latest version of Firefox and install it on a thin client as you can with a workstation. The manufacturer must provide a special image for your make and model. This is something that needs to change, but for now, the software that the manufacturer makes available is a crucial factor in selecting a thin client. If you want Firefox on a standalone

thin client, the manufacturer has to provide it. If you want Flash and Java to work, the manufacturer must provide the plugins. Don't expect the plugins to be current releases either. The size of some plugins has outpaced even the plummeting cost of memory. In particular, Acrobat and Java have grown so enormous that it is more reasonable to use an older release than pay for the additional Flash and RAM required to run them.

How software is made available depends on the manufacturer. There are basically two methods. One is to provide individual modules. This allows you to pick and choose, but more labor is involved in preparing the clients. The other method is for the manufacturer to provide monolithic images with all the options needed. This can be practical if the manufacturer is flexible about providing custom images.

When using thin clients in client mode, the applications are all normal installations on the terminal server, which is simply a high-performance server with enough horsepower to do the application processing.

In client mode, the thin client has a dual nature. It is a client in respect to the application services provided by the terminal server, but it is also a server in respect to providing those applications with access to local hardware. The local hardware being served up is primarily a keyboard, video and mouse (KVM), but there also can be local audio, USB storage devices and printers.

Thin clients are available with Linux, Windows CE and Windows XP Embedded. Barring some desire to use Internet Explorer in standalone mode, there really isn't any reason to consider anything but Linux for a thin client. Even if the terminal server is MS Windows, the fact that Linux is running on the thin client is completely transparent to the user. CE and XP only add software license costs to each client, and XP doubles the Flash and RAM memory requirements on the client (128MB minimum Flash and RAM for Linux vs. 265 Flash and RAM for XP). Because of this, the most commonly deployed thin-client configuration today is Linux thin clients connecting to MS Windows terminal servers.

### Thin-Client Protocols

There are four common thin-client protocols:

- Remote Desktop Protocol (RDP) is a proprietary MS protocol that provides monolithic remote desktop support. It includes encryption and redirection to allow remote applications to access most local hardware, including audio, filesystems and printers. It currently does not allow single applications to be run remotely (without a desktop), but RDP 6.0 is supposed to add this. RDP clients are available for Linux, but there is no functional RDP server, although a nascent product named xrdp is under development.
- Independent Computing Architecture (ICA) is a proprietary protocol from Citrix. It is largely similar to RDP, which is based on an earlier version of ICA. ICA includes the ability to run single applications remotely, without the entire desktop, but it requires Citrix Presentation Server, which is available for MS Windows and some UNIX systems.
- X Display Manager Control Protocol (XDMCP) is an open standard used by the X Window System (X). It is notably different from RDP and ICA in two respects. First, the same

software modules (described below) are used for local and remote sessions. No separate terminal server software is necessary. Second, it is not monolithic. In the UNIX tradition, it does what it does and works with other tools that do what they do. It does not provide compression or the ability for remote applications to access local hardware other than KVM.

- NX is an open standard server built on top of X that simplifies thin-client networking. It includes built-in support for encryption (using SSH), access to the local filesystem (using Samba) and local audio (using ESD or aRts). The server also is able to translate foreign protocols to allow connections from RDP and other clients. NX is a product of NoMachine, which develops an open-source core, on which proprietary versions of both the server and client are built. There is also a completely open-source project called FreeNX.

A distinction should be made between these protocols and remote framebuffer protocols like VNC. VNC provides remote control of a desktop that is still local, while thin-client protocols provide remote desktops.

### Pieces of X

X is nothing if not modular. Modularity is a good thing, but seeing how all the pieces of X fit together can be daunting for a new user. Below is a summary of the main modules and their interactions that will make the rest of this article accessible to readers with no previous

## DEDICATED SERVERS ROCK BOTTOM PRICING



- » 1300GB/mo Bandwidth
- » 1GB RAM DDR400!
- » 160 GB SATA2 Hard Drive
- » 100Mbps Dedicated Port
- » FREE REBOOTS 24x7
- » FREE SUPPORT 24x7
- » FREE SETUP Available\*
- » BRAND NAME Hardware

\$**60**  
/mo  
Celeron D

\$**160**  
/mo  
Dual Core Xeon

**w.cari.net**

\* See website for details.

**carinet**  
betterServers. betterService.

X networking experience.

All access to the physical display is through the X server. This is a source of confusion for new users, because the display is intuitively client-side. But, the display is the service to which it provides access, hence the name. The clients for an X server are X applications that use it to display their output. We will see relationships later in this article where the X server acts as a client to other services.

The display manager (DM) is the heart of the terminal server. X servers and DMs have a dual client-server/server-client relationship. An X server can, as a client, initiate a connection to a DM on UDP port 177. The DM will then connect to the X server as a client on TCP port 6000 to display a graphical login screen to the user. A client can have multiple displays (windows or virtual terminals), in which case, the second display would be on port 6001 and so on. The protocol for this communication is XDMCP. If the X server and the DM are on the same system, they communicate using a UNIX socket.

The X server and DM are about displays and pixels. Neither has any concept of a window or a widget. These are handled by the window manager (WM).

Although the window manager provides the fundamental functionality and the major aspects of the look and feel, that is not enough to consider it a fully usable system. The desktop environment (DE) completes the user interface with utilities, such as control panels and toolbars, and basic applications, such as calculators and text editors.

There is often one additional component used: an X font server. The name of this server is `xfs`. In relation to `xfs`, an X server is a client that connects to an `xfs` server on TCP port 7100. X servers also can be configured to retrieve fonts from a filesystem folder.

The main decisions to make when deploying thin clients and a terminal server are the DM and the DE. The X server is built in to the thin client, and the DE will have a default WM that there is usually no reason to change. There are two dominant DEs in use today: GNOME (GNU Network Object Model Environment) and KDE (K Desktop Environment). Both have extensive features, and they are about equal in market share. GNOME is written in C and uses the GTK+ libraries. KDE is written in C++ and uses the Qt libraries. Both GNOME and KDE have their own WMs, named Metacity and KWin, respectively. They each also provide their own DMs, GDM and KDM, one of which is normally used in place of the standard XDM provided with X.

## Terminal Server Configuration

Start by installing your distribution of choice. The specific file locations given below are for Fedora 5. Most distributions install only one DE by default, so make sure to select the desired DE during installation. Although many distributions, including Fedora, give the choice of GNOME or KDE, some have opted to provide only one DE.

GNOME and KDE coexist well. One is set as the system default, but both GDM and KDM allow you to select GNOME or KDE desktops on the fly for each login. The system default DM, on the other hand, is the only DM that will be used.

To select the system default DE and DM, edit `/etc/sysconfig/desktop`. It should have only two lines:

```
DESKTOP="DE"
DISPLAYMANAGER="DM"
```

DE is either GNOME or KDE, and DM is either XDM, GNOME or KDE.

If the system does not automatically boot to a graphical login, change the default runlevel (`initdefault`) to 5 in `/etc/inittab`.

To use a font server, run `ntsysv` and select `xfs` to run at boot. Also, configure `xfs` by editing `/etc/X11/fs/config` and remove the line `no-listen = tcp` to allow outside connections to `xfs`.

KDM is configured using the file `/etc/kde/kdm/kdmrc`, which is in

INI format. To allow remote connections, set `Enable=true` in the `Xdmcp` section. You probably also will want to customize the `X*-Greeter` section, which controls the appearance of the login screen. Note that if the `UseTheme` parameter is true, many other parameters in this section will be overridden. KDM also can be configured using the KDE Control Panel, but it loses all the comments in `kdmrc`. I prefer to edit `kdmrc` directly.

GDM also has an INI format configuration file (`/etc/gdm/custom.conf`). Simply set `Enable=true` in the `xdmcp` section. The GDM configuration file is not heavily annotated, so the GUI configuration tool `gdmsetup` may be preferable. Run `gdmsetup` locally on the terminal server. On the Remote tab, change `Style` to Plain, Plain with face browser or Same as local. If in doubt, choose Plain.

For any DM, access is controlled by the file `/etc/X11/xdm/Xaccess`. Simply add the IP address or DNS name of each allowed host. A `*` on a line by itself will allow connections from any host.

This is everything necessary to allow a thin client to log in to a desktop on the terminal server, but more server configuration will be necessary later to access local thin-client hardware beyond KVM.

## Connecting with X

A thin client is not the only way to access the XDMCP server. Client software also can be run from a workstation. You can access either a desktop or directly run applications.

If X is not currently running, the following command provides a login to a remote desktop on a terminal server host:

```
X -query host
```

If X is already running, the same command also will work, with the local desktop and the remote desktops being on separate virtual terminals (VTs). To open the remote desktop in a window, use:

```
Xnest -query host
```

If either of these give the error "Server is already active for display 0", select a different display number by adding `:1` as the first option.

To run an X application remotely without a desktop, use:

```
ssh -X -l username host
```

to log in, and then run the application from the command line. The `ssh` option `-C` will add compression for slow connections.

If your workstation has `tsclient` installed, this provides a GUI front end for `Xnest` as well as client software for other protocols.

Most of the above functionality can be added to an MS Windows workstation by installing `Cygwin`. When installing `Cygwin`, add the package `X11/xorg-x11-xwin` to the default installation. Also, add `Net/openssh` if you want to use SSH. The above commands should then work in a `Cygwin` shell. For SSH, run `startx` first, then run `ssh` in the X terminal window that it creates.

## Thin Client to a Terminal Server

Thin clients provide GUI tools that make basic network configuration similar to any other network appliance. Beyond the basic configuration, these tools allow the creation of sessions. A session defines the server to access and the protocol to use. XDMCP can be a session, but because the thin client is running X natively, connecting to an XDMCP server might be a basic display configuration setting. How this is done depends on the manufacturer. Other protocols are always configured as sessions.

The configuration details below were tested on an Igel 364 LX. They should be general enough to work with any thin client with ESD and NFS server capabilities, but these are not features that should be assumed.



## Local Audio

Many thin clients have no support whatsoever for local audio from a Linux terminal server. Those that do typically have only ESD. This requires that the applications be configured to use ESD (most have this option, but not all). The following also must be added to the `.bash_profile` of thin-client users to identify the IP:port of the thin client's ESD server:

```
export ESPEAKER=${DISPLAY%:*}:16001
```

## Accessing Local Storage

Because thin clients have no built-in drives, the only local storage of interest is USB-connected. We want locally inserted devices to be accessible from a desktop icon. But as the desktop is running on the terminal server, we need to make the terminal server see these local files.

This requires a thin client with a local NFS server configured to automatically detect and share USB devices. On the terminal server, we configure the `autofs` daemon to detect these remotely mounted devices automatically and mount them locally. Create a directory `/etc/auto` on the terminal server. For each user that is allowed to access local storage, create a file `/etc/auto/username` with the following contents:

```
usb -rw,soft,intr 192.168.0.64:/autofs/usb0
```

Replace 192.168.0.64 with the thin client's IP address, and the path `/autofs/usb0` will vary by manufacturer. Create a directory `/home/username/media`, then add the following to `/etc/auto.master`:

```
/home/username/media /etc/auto/username --timeout=15
```

Finally, create a symlink on `username`'s desktop to `/home/username/media/usb`. The user now can insert a USB drive, and clicking the symlink will cause `autofs` to mount it on the terminal server.

This method works and has been used in real deployments, but it has an inherent limitation. The thin clients must have static IPs, and each user is tied to an IP address. In cases where users need to float between stations, this will not be adequate.

## Restricting Physical Login Locations

In many cases, it is actually required that user access be restricted to specific locations. This is easily accomplished using the PAM login access control table. First, the thin client must be given a static IP address. Then, add the following entry to `/etc/security/access.conf` on the terminal server:

```
-:username:ALL EXCEPT 192.168.0.64
```

The format of this file is `permissions:users:origins`. So the above example removes (-) permission for user `username` from all addresses except 192.168.0.64.

Besides the obvious security application, this is also useful for public-access thin clients. While creating a separate generic account for each thin client (`user1`, `user2` and so on) gives each one a separate home directory so users will not trip over each other, it is easy to log in accidentally using the wrong generic account at a given workstation. This procedure prevents that.

## Conclusions

Thin clients have matured and are ready for widespread use. Their benefits are too compelling to ignore, and most have a commitment to Linux as their primary platform. Unfortunately, most are myopically focused on MS Windows terminal servers and are neglecting support

for Linux on the server side. As they become more widely deployed, the ironic possibility of Linux systems becoming an impediment to the deployment of open source on the desktop is very real.

Some specific items that must be addressed are:

- Thin clients are too proprietary. Open tools are needed for building Flash images and other system management tasks.
- Universal support for full-duplex, low-latency audio.
- Secure, easy and mobile access to local USB storage devices.
- Support for local non-PostScript printers.
- Encryption and compression.

The solution is likely NX or something very similar—something that retains the modularity of the system while integrating the components into a cohesive whole. I have not yet seen a thin client with a fully functional NX client. ■

**Resources for this article:** [www.linuxjournal.com/article/9388](http://www.linuxjournal.com/article/9388).

Lyle Frost is a consultant with Citadel Network ([www.citadelnetwork.com](http://www.citadelnetwork.com)), an IT management firm in Indiana.

<b>LAYER 42™</b>	
● Redundant UPS and generator	
● Nationwide network	
● Free tech support	
<b>2U</b> 256kbps ~80GB \$60/mo.	<b>4U or Mid-tower</b> 256kbps ~80GB \$80/mo.
<b>1/4 Rack</b> 512kbps (14U) ~165GB \$200/mo.	<b>1/2 Rack</b> 1mbps (28U) ~330GB \$350/mo.
<b>www.layer42.net</b>	
All prices include 100Mbps port, Firewall, 24x7 Monitoring and DNS hosting	
408-450-5740	2312 Walsh Ave., Santa Clara, CA 95051

# AcidRip—a Gtk2 Front End to MEncoder

How to use AcidRip to make DVD backups. DANIEL BARTHOLOMEW

**MEncoder** is a wonderful little command-line utility included with MPlayer for encoding video. It can take as its source file any video format that MPlayer can read, including Windows Media, MPEG-2 (DVD), QuickTime, MPEG-4, DivX and many others. It then can convert those source files using several encoders, such as lavc, libdv, xvid and x264.

The reasons for doing what basically amounts to conversion from one digital format to another format—possibly the exact same—are several. Converting from the North American NTSC standard framerate of approximately 29.97 frames per second to the European PAL standard of 25 frames per second is one reason. Removing dust and scratches and performing color correction are others. My reason is disk space—newer video codecs like xvid and x264 can do more in less space than older formats, such as MPEG-2. With no detectable loss in quality, a 4GB DVD movie easily can fit in much less than 2GB of space. Furthermore, if you are more aggressive and don't mind scaling the picture, you can shrink it further so that it will fit on a CD-ROM. Even at that size, the picture and sound quality can still be excellent—if you know how to use MEncoder.

Sadly though, MEncoder is not that easy to learn to use properly. The man page alone clocks in at 7,216 lines. Given time and patience, I am sure it is possible to learn the ins and outs of this wonderful program, but I do not have much patience, and I have no time.

The problem is this: my children seem determined to break every DVD in the house. It's not that they are trying to, they're just being kids, but children and DVDs are a bad mix. DVDs are simply too fragile. They seem to get scratches and cracks as soon as you open the case the first time. My little angels have already destroyed *Shrek*, *Ice Age*, *Black Beauty* and *Chitty Chitty Bang Bang*, among others—I would prefer the destruction to stop there.

My plan is to back up every DVD in the house onto my Linux server. Then, using MythTV or another suitable front end, enable the kids to watch their movies on

the television as much as they like. The original DVDs, meanwhile, will be placed carefully in their cases and locked away where little fingers cannot get to them.

There is a lot of storage on my server, but at four-plus gigabytes per disk, and with a growing library of DVDs that is already more than a 100 discs, I don't have that much storage. This is where MEncoder comes in—sort of. I need to convert my DVD library from MPEG-2 into a more storage-friendly format, and MEncoder can do it, but it has me beat—at least for now.

The design goal of MEncoder seems to be to give you the ability to tweak every aspect of your encoding, from format to framerate to bitrate to dimensions to color. With this much power, I have found it very easy to make many errors. Others also have gone through the struggle to learn MEncoder, and thankfully, some of them have tried to make it easier to use. The result is not perfect, but it is a step in the right direction.

AcidRip is a Gtk2::Perl front end to MEncoder. It guides you through setting the options for MEncoder and warns you if you try to do something that will result in a less-than-stellar outcome.

You can download AcidRip from the SourceForge product page. It is also in the package repositories of some Linux distributions. Because AcidRip is a Perl program, once you have unpacked the source files, you can launch AcidRip right from the source folder.

AcidRip depends on MPlayer and MEncoder, so you need to have them installed and working. You also need the DeCSS package to enable the reading of encrypted DVDs. Basically, if you can use

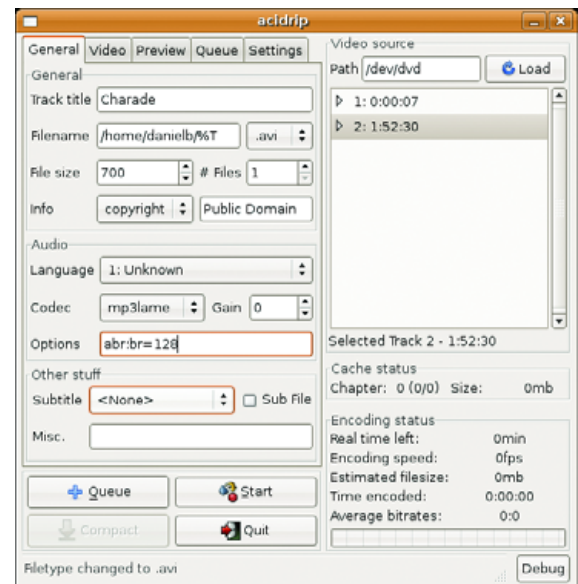


Figure 1. The AcidRip General Settings Tab

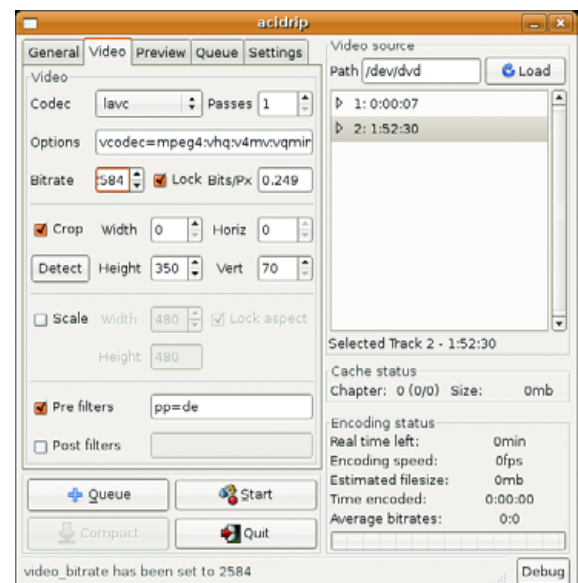


Figure 2. The AcidRip Video Settings Tab

MPlayer to watch a DVD, you can use MEncoder to rip it. MPlayer and MEncoder are included with most distributions, so there more than likely is a prebuilt package

available in your distribution's package repositories. If not, download the MPlayer source and essential codecs packages from the MPlayer Web site. Follow the installation instructions (see the on-line Resources), and you should be in business.

The AcidRip program also utilizes a little program called `lsdvd` that is included in the AcidRip source package.

After starting AcidRip, the first thing you need to do is load a DVD. Load the DVD in your DVD drive, and press the Load button in the Video Source section. You will see a listing of all of the chapters and tracks on the DVD. My example DVD has only two tracks; other DVDs could have one or several. To view the chapters in each track, click on the disclosure triangle.

Look for the longest track; this will be the actual movie. If you are ripping a DVD filled with behind-the-scenes extras, interviews and such, there may be many short tracks with no real indication of which one is the main movie, so you may have to try a few until you get to the one you want to encode. Select the track you want to encode by clicking on it.

Now that we've selected the track to encode, we need to set a few options. First, under the General tab, put in the track title. This will end up being the file-name for the resulting .avi file. In the Filename field, put in the path to where you want to save the file, ending with %T (the default save location is your home directory), but you can set it to wherever you want. We will get to the file size and number of files boxes later.

If you like, you can add some basic meta-data about the movie you are ripping into the Info box, such as the name, artist, subject, genre and copyright information. This information can be read by MPlayer on playback, but otherwise, it is not very useful.

In the Audio section, you can leave the selection on "<Default> English" or choose another audio track using the drop-down menu. Be careful when selecting audio tracks, as some may be commentary tracks, and on some DVDs, certain entries may simply be blank.

From the Audio Codec drop-down box, choose how you want to encode your audio. The choices are dependent on the codecs you have installed. On my machine, they are copy, pcm, mp3lame, lavc and faac. For speed, copy is the fastest, as it simply copies the audio track from the DVD straight into the resulting .avi file. If you choose to encode your audio as MP3 using the mp3lame or lavc codecs, you can adjust the bitrate, but the higher you set the bitrate, the longer your encoding will

## MANUAL CROPPING

Some DVDs may give AcidRip trouble due to fuzzy borders, so if it returns with a "crop failed" message, you may have to set the crop manually. If this is the case, start with a width of 720, a height of 480 and go down from there. 720x480 is the size of a standard NTSC DVD frame. For PAL DVDs, the size is 720x568.

When cropping manually, the Horiz and Vert sections can be a little confusing. What they are is the offset from the top-left corner of the full frame of where the crop frame should be positioned. The crop frame itself is specified in the width and height boxes.

For example, take the image shown in Figure I from the film *Charade*, starring Audrey Hepburn and Cary Grant. This film is good to use as an example for two reasons. First, AcidRip could not detect the the proper crop settings. Second, due to a quirk in United States copyright law, when this film was released, *Charade* fell into the public domain, and it can be used by anyone for any purpose, including this one.

As you can see from Figure I, there is a black bar down the left side of the film, which exists throughout the film. There is also a small black border down the right side. Both borders are fuzzy, and therefore they are hard to detect accurately. The top and bottom borders of the film also exhibit some variation throughout. The manual crop settings that eventually worked for me were:

Width: 705  
Height: 346  
Horiz: 11  
Vert: 71

This is simply another way of saying that my crop rectangle is 705 pixels wide and 346 pixels tall. The crop rectangle is offset from the left edge by 11 pixels and from the top edge by 71 pixels.

To get at these final dimensions, I first started with a width of 700, a height of 400, a horizontal offset of 10 and a vertical offset of 60, which were my best guesses of what the dimensions were. Then, by switching back and forth between the Video and Preview tabs, I was able to fine-tune the settings until I was happy with the result.




Figure I. An Example Frame from the Movie *Charade*

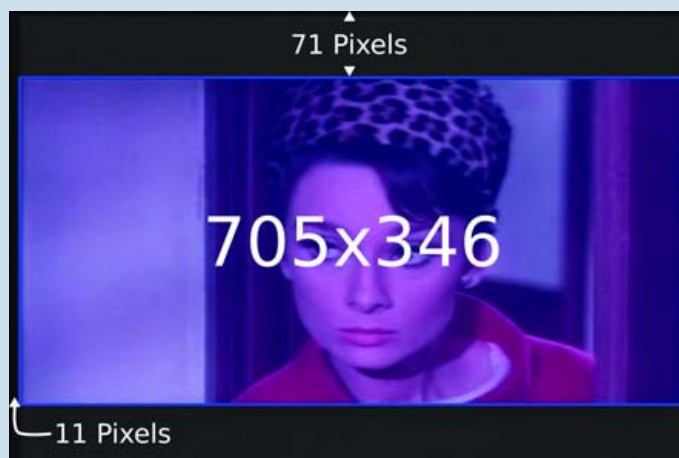


Figure II. The final crop settings I used with this movie.

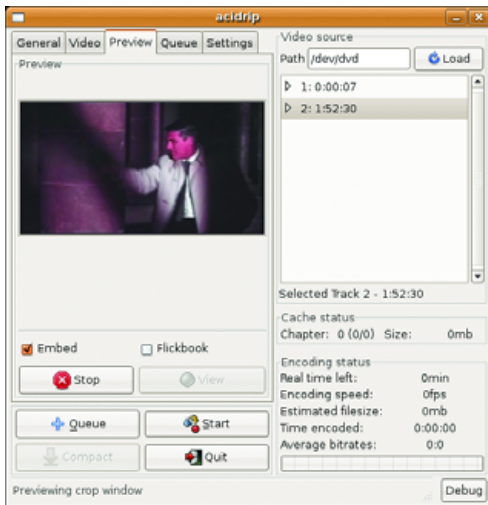


Figure 3. Preview your settings here. If you don't see a picture or hear any audio, you probably need to tweak your settings.

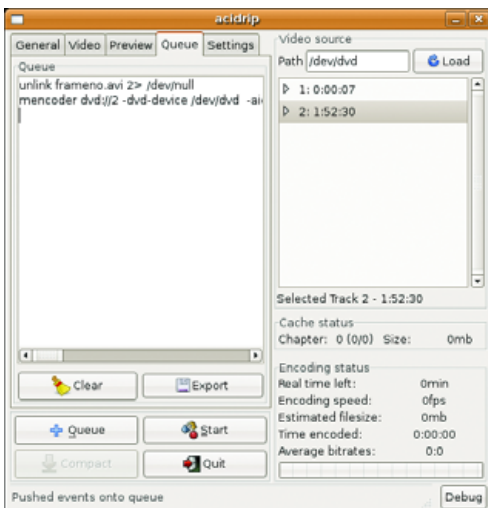


Figure 4. The Queue tab—this shows you the MEncoder command that AcidRip will run to rip your DVD.

take. If you find your encodes are too soft or too loud, you can adjust the gain up or down as needed. I have not found this to be necessary in most, if not all, cases.

Now we need to set the video options. Click on the Video tab to view them. Again, the codec choices you see are dependent on which codecs you have installed. On my machine, the choices are copy, raw, nuv, lavc, vfw, qtvideo, libdv, xvid and x264. For the best quality in relation to file size, use the x264 codec. I also have had great success with the lavc codec set to its default values. We'll get to the Passes, Bitrate and Bits/Px boxes later.

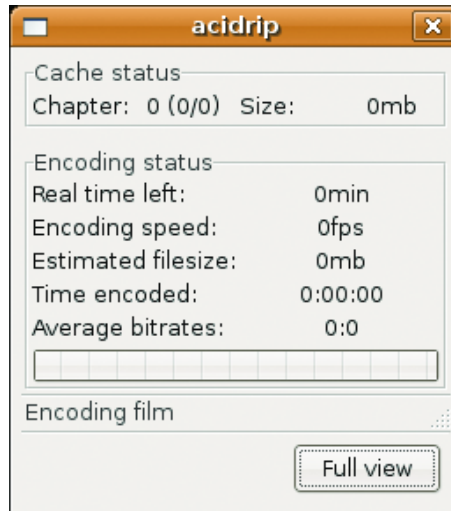


Figure 5. The Incredibly Uninformative Progress Dialog

It is always a good idea to select the crop check box, especially when ripping a widescreen DVD. The last thing you want to do is waste a lot of time and disk space encoding the two black bars at the top and bottom of the video frame. By pressing the Detect button, AcidRip uses MPlayer to skip around to several different individual frames of your selected video track. MPlayer then uses those frames to try to guess at the appropriate crop settings. You can fiddle with the Width, Height, Horiz and Vert to adjust the auto-detected crop if you want, but I usually leave them as they are. For manual cropping, see the Manual Cropping sidebar.

If you are trying to fit the entire film on a single CD-ROM, you may want to scale the picture. The scale feature scales after the cropping is done, so don't try to adjust the crop to fit the scale, simply enter in the scale size you want. Also keep the Lock aspect check box ticked to avoid a distorted picture.

The final options on the Video tab are to adjust the Pre and Post filters. I usually leave them alone.

If you choose lavc as your video codec, you can fiddle with the Bitrate, Bits/Px and set the number of encoding passes you want to use. Generally speaking, the optimal Bits/Px setting is right around 0.249 for MPEG-4 video. If you tick the Lock check box you can adjust the Bitrate manually until you arrive at a Bits/Px setting of right around 0.249. Multiple passes can and will greatly increase the encoding time, but they also will increase the file quality.

By locking the bitrate, you will have no control over the size of the resulting file.

## Advanced Encoding

Now that I am familiar with encoding with AcidRip, my next project is to try to take it to the next level by using MEncoder directly. AcidRip's queue export feature really helps with this. You can export a small shell script of the exact commands that AcidRip passes to MEncoder. Using that as a starting point, I can tweak the settings even further.

The MEncoder documentation is also a great source for encoding instructions and ways to tweak the parameters to get the best image quality. Now if I could only find the time.

After setting your video options, switch back to the General tab, and you will see that an estimated file size has been entered. If you would like to determine the file size manually and have the bitrate adjusted accordingly, untick the Lock check box on the Video tab, switch back to the General tab and adjust the file size. This is normally done when you want to fit a DVD onto one or more CD-Rs.

If you are planning to burn the ripped DVD onto CD, there are a couple of ways to go about doing it. First is to encode at full size into multiple files. Four 700MB CDs are usually enough to hold a full-size movie at good quality. The second option is to set your target file size and then scale down until the Bits/Px is good.

When adjusting the scale size, use the up and down arrows. That way, you will be able to see the bitrate adjust in real time. There is a glitch that prevents auto-updating of the other fields if you type them in manually. This applies to all fields that you can adjust.

Once you think you have your settings correct, switch over to the Preview tab. Keep the Embed check box ticked and the Flipbook check box unticked, and click the Preview button. As long as there are no errors in your settings, the movie will play. You may find that you have to adjust your crop settings. One thing I like to do to get a better idea of what the movie will look like is to select a chapter from the middle of the movie in the Video Source section instead of watching the movie from the beginning. When doing this, I just have to be sure to change it back before encoding

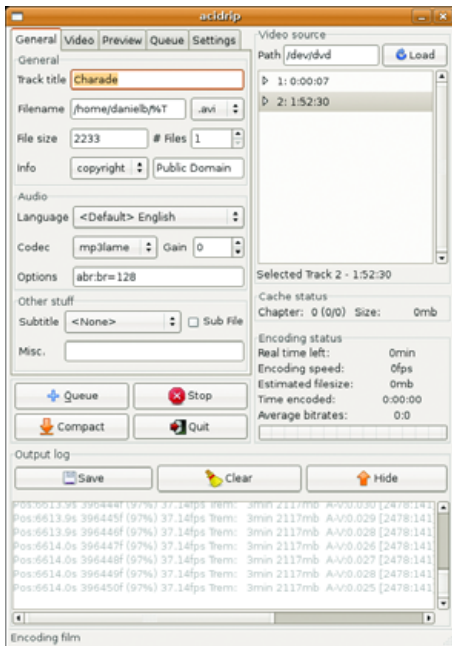


Figure 6. The Incredibly Informative Output Log

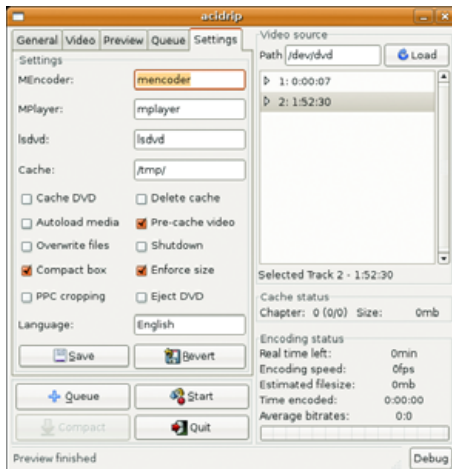


Figure 7. The AcidRip Settings Tab

for real. When you have seen enough, press the Stop button.

Once you are happy with the settings, you are ready to queue the film for encoding by pressing the Queue button. Switching over to the Queue tab, you can choose to clear the queue or to export the current queue as a shell script. You also can set up multiple encodes to run sequentially, which is very useful for encoding a group of behind-the-scenes extras all into their own files.

Once your queue is set up with the encode or encodes you want, press the Start button to begin. You then will be

presented with a small progress window showing—nothing. This is probably the biggest glitch of AcidRip. It's possible that this might be fixed by the time you read this, but with the current versions of AcidRip (0.14) and MEncoder (1.0pre8), the display is broken. MEncoder is working though; AcidRip is just not telling you about the progress.

To view the progress, click on the Full view button to return to the regular interface and then on the Debug button to view MEncoder's raw output. Scroll to the bottom, and you will see its progress. You should also view the debug window if AcidRip fails to encode the movie for some reason, as it can provide you with good clues as to why and which option caused the encoding to fail.

The final tab in AcidRip is the Settings tab. There you can tweak various settings, including the paths to the MEncoder and MPlayer applications—useful if, for example, you have them installed in nonstandard

places. You also can set other options, which are fairly self-explanatory.

In conclusion, AcidRip is a very useful application that, for me at least, helped me get a handle on ripping my DVD collection to my computer. It could use some bug fixes to correct the interface glitches, but apart from those, it works and works well.

The only really unfortunate thing about AcidRip is that the author, Chris Phillips, has stated that he is not interested in updating the product much, if at all. But, due to the beauty of open source, an energetic Perl hacker easily could fork the project and make the necessary updates. Any takers? ■

**Resources for this article:**  
[www.linuxjournal.com/article/9389](http://www.linuxjournal.com/article/9389)

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.

**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](http://www.ironsystems.com)

**CALL: 1-800-921-IRON**

# Federated Desktop and File Server Search with libferris

How to federate CLucene personal document indexes with PostgreSQL/TSearch2. BEN MARTIN

The **libferris project** has two major goals: mounting anything as a filesystem and providing index and search for anything it can mount. Using libferris to provide desktop search was described in my February 2005 article, “Filesystem Indexing with libferris” in *Linux Journal*. The indexing capabilities of libferris have grown since then. One new feature is to allow a group of indexes to function logically as a single, “federated” index. This lets you have an index for your file server, another for your man pages and a third for your personal documents. You then can run queries against all three as though they were a single index.

libferris handles its index and search using a plugin system. There currently are index plugins for db4, PostgreSQL, ODBC, Redland (RDF), Xapian, Beagle, Yahoo, LDAP, CLucene, Lucene and external processes. The indexes that form a federated index in libferris can use any mixture of those index plugins.

libferris has two different types of indexing plugins: full text and metadata. The metadata interface of libferris is based on the Extended Attribute (EA) kernel interface. Having two index plugin types allows the index plugin to organize data on disk to best support queries.

A full-text index normally will maintain for each word from a human language a list of which files contain that word and a statistical measure of how important that word seems to the document. The statistic allows documents that are “more relevant” to be presented first in the results. Such statistics normally relate to how large a file is, how often the word appears in that file and how rare the occurrence of that word is across all indexed files.

A metadata index has to associate a docid with a keyword and value. For example, /tmp/foo has a size of 145. The metadata index has to be able to process queries, such as `size>=4kb && modified this week`, and efficiently return the docids for files that satisfy this query. The main difference between metadata and full-text index plugins is that the metadata queries contain value comparisons on metadata (for example, `mtime>=last week`), whereas full-text queries generally are more interested in the presence of a word in a file.

## The User View

From an index user’s point of view, having this distinction is an annoying implementation artifact. To get around this, a full-text index can be linked to a metadata index using the `feaindex-attach-fulltext-index` command. Queries combining both metadata and full-text searching can then be executed on the metadata index. It is convenient to think of the metadata index as owning the full-text index.

The metadata query format reserves any metadata names starting with `ferris-` to have special meaning. A metadata name `ferris-fulltext-query` or `ferris-ftx` will execute its query value as a full-text query on the linked full-text index. Shown in Listing 1 is a metadata query seeking all files under a given size with the two given words in them. If instead of combining the results with `&`, we used the `or` operator `|` in the query, any results matching either subquery would be returned. To query a full-text index, the `findquery` command is used. Combined metadata and full-text indexes are queried using the metadata query command `feaindexquery`.

Listing 1.

### A Combined Full-Text and Metadata Index Query

```
$ feaindexquery \  
'(&(size<=250k)(ferris-ftx==alice wonderland))'
```

The above discussion of docids becomes relevant when combining two types of index plugins like this. The greatest efficiency can be gained when both the metadata and full-text index plugins are using the same storage—for example, the PostgreSQL (metadata) and TSearch2 (full-text) plugins using the same underlying PostgreSQL database, or both indexes using the same CLucene storage.

The efficiency is obtained because each URL has the same docid. Using the PostgreSQL combination as an example, to resolve the query from Listing 1, the full-text subquery will be run against the TSearch2 plugin obtaining a set of matching docids. The set of docids matching the size query is obtained, and the set intersection of the size and full-text query results is returned. This final step can be done only if it is known that both the metadata and full-text index have the same docid for the same URL. Otherwise, the docids from the full-text query have to be converted into URL strings and then into the docids of the metadata index first.

When using a metadata and full-text plugin together like this, make sure that each file is added to both indexes.

Each metadata index plugin will automatically detect if it is safe to use the docids of the full-text index directly that is linked to it.

The federation index plugin is a metadata plugin. A federation is formed using many metadata indexes with one nominated as the writable index. As each metadata index can own a full-text index, this allows federations of an arbitrary number of full-text and metadata indexes. Each index in the federation can be updated independently of the federation.

## Setting Up a CLucene and PostgreSQL Federation

Indexes are created using either the `fcreate` or `gfcreate` tools. The former is a command-line tool, and the latter has a GTK+ 2 GUI. In this article, I use the `fcreate` command. To find out what other options are available during index creation, simply replace `fcreate` with `gfcreate`, and a GUI will be presented. Both metadata and full-text indexes reside in a directory, even if only configuration settings are saved in that directory. For example, using the PostgreSQL plugin, the indexed data will be in a PostgreSQL database and only a small config file will live in the filesystem directory. Using directories like this allows you to tell libferris which index to use by passing a filesystem path.

Some shell scripts are distributed with libferris to help set up indexing. For PostgreSQL and CLucene, these scripts start with `ferris-recreate-primary-fulltext-and-eaindex-as` and end with either `clucene` or `postgresql`. Both are geared to set up your default metadata and full-text indexes using the nominated index plugin.

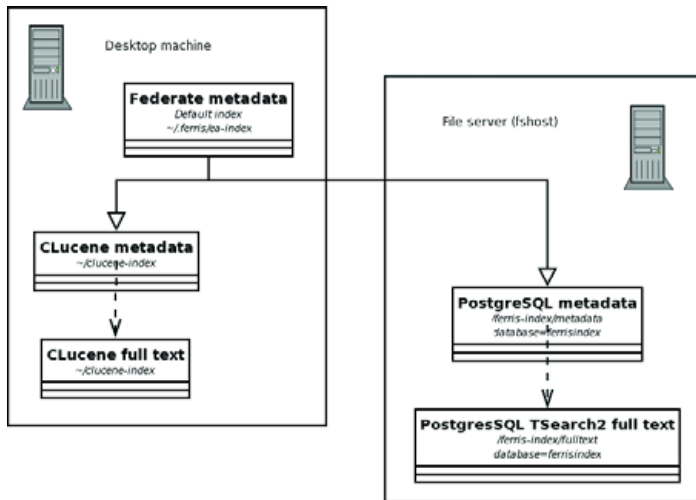


Figure 1. The Federation of Indexes

Your default indexes are stored in subdirectories of `~/ferris`.

We'll make our default index a federation of a local CLucene index for personal files and PostgreSQL for a file server. This means we will have five indexes in total: the federate metadata index, a metadata and full-text CLucene index, and a metadata and full-text PostgreSQL index.

The two CLucene indexes will be linked together, and the two PostgreSQL indexes will be linked to each other. We can use the default path in `~/ferris` for the federation index. We will put the CLucene indexes in `~/clucene-index`. I'll assume the machine that will run PostgreSQL and maintain the file server index is a server called `fshost`. The index can be on a different machine from the actual file server if desired. The contents of many file server machines and other documents can be added to the file server index if you like.

For PostgreSQL indexes, the directory for the index will have only a configuration file in it. This file will contain information telling the index plugin where the database is located and what user name and password to use to connect. I'll assume we are creating the PostgreSQL file server indexes in `/ferris-index` on the file server, though any path is fine. To make things simple for people who are intended to use this index, having its directory on the file server makes its use in a federation simple. We'll use the PostgreSQL database name `ferrisindex`. The setup is shown in Figure 1.

To use CLucene for local indexing, we can use the `clucene` recreate script with a minor modification for the index paths, as shown in Listing 2. Notice that the second `fcreate` has the `db-exists=1` parameter to tell the index plugin that there is an existing CLucene

#### Listing 2.

### Setting Up Two CLucene Indexes

```
$ mkdir -p ~/.clucene-index
$ cd ~/.clucene-index
$ fcreate `pwd` \
  --create-type=fulltextindexclucene
$ fcreate `pwd` \
  --create-type=eaindexclucene db-exists=1
$ feaindex-attach-fulltext-index \
  --ea-index-path `pwd` \
  --fulltext-index-path `pwd`
```

index at this path. This places both metadata and full-text information into the same CLucene index.

Make sure that metadata you want to use in queries is not listed in `attributes-not-to-index` and will not match `attributes-not-to-index-regex` for the index. Run `gfcreate /tmp --create-type=eaindexclucene` to find your current default values for these parameters.

Setting up a PostgreSQL/TSearch2 combination is a two-step process. The first step, using the `ferris-setup-template-index-database.sh` script, creates some template databases and needs to be done only once. The script assumes it is being run on the host that has the PostgreSQL database on it. This script installs Generalized Index Search Trees, TSearch2 and PL/pgSQL into two template databases that the metadata and full-text plugins take advantage of. Some of these features live in a `postgresql-contrib` package in many distributions.

The commands shown in Listing 3 create a TSearch2 full-text index and a metadata index in the same database on host `fshost`. These will

#### Listing 3.

### Commands to Run on the File Server to Create PostgreSQL Indexes

```
$ ferris-setup-template-index-database.sh
$ mkdir -p /ferris-index/metadata
$ mkdir -p /ferris-index/fulltext
$ cd /ferris-index
$ fcreate /ferris-index/fulltext \
  --create-type=fulltextindexsearch2 \
  dbname=ferrisindex host=fshost
$ fcreate metadata \
  --create-type=eaindexpostgresql \
  host=fshost dbname=ferrisindex db-exists=1
$ feaindex-attach-fulltext-index \
  --ea-index-path metadata \
  --fulltext-index-path fulltext
```

#### Listing 4.

### Re-Creating Default Indexes Using PostgreSQL

```
$ mount fshost:/ferris-index /ferris-index
$ fcreate ~/.ferris/ea-index \
  --create-type=eaindexfederation \
  primary-write-index-url=~/.clucene-index \
  read-only-federates=\
  "~/clucene-index,/ferris-index/metadata"
```

#### Listing 5.

### Examine Index Metadata and Change the User Name

```
$ cd /ferris-index/metadata
$ ferrisls -lh ea-index-config.db
11  cfg-idx-database
6   cfg-idx-host
...
$ fcat ea-index-config.db/cfg-idx-host
fshost
$ echo -n foouser | ferris-redirect \
  --trunc ea-index-config.db/cfg-idx-user
```

reside in /ferris-index as mentioned before. This directory should be readable over the network by those who are intended to use the index. Below, I assume this is exported using NFS and access the path using fshost:/ferris-index. These indexes are then linked together to allow combined queries. Make sure that the db files in /ferris-index are readable by those who should be able to access this index.

Back on the desktop machine, we then create a federated index combining the local CLucene and remote PostgreSQL indexes, as shown in Listing 4.

This assumes that the parameters used to create the PostgreSQL indexes are valid for the desktop user. As libferris knows how to mount db4 files, changes to the configuration settings can be done with libferris clients. See Listing 5, which uses the ferris-redirect command to allow shell redirection into any libferris file.

The federation index plugin delegates all of its work to other existing indexes. Because of this, we nominate that when files are added to the federate index, then the federate plugin should delegate the add to the CLucene plugin maintaining our personal index.

### Populating Indexes

Most index plugins will detect whether a file has not changed since it was indexed and automatically skip it upon re-indexing. At least the Xapian, Redland, CLucene and PostgreSQL plugins support this. Those plugins that do not currently support this will issue a warning. This allows a cron job simply to run find to list files that should be in the index and pipe them to feaindexadd.

Shown in Listing 6 are commands to populate both indexes. Note that when using CLucene for both full-text and metadata indexes in a

Listing 6.

#### Adding Files to an Index

```
# Local index
$ find ~ -name ".*" -prune -o -print | findindexadd \
  -P ~/clucene-index --filelist-stdin
$ find ~ -name ".*" -prune -o -print | feaindexadd \
  -P ~/clucene-index --filelist-stdin

# File server index, run on fshost
$ find /documents | findindexadd \
  -P /ferris-index/fulltext \
  --filelist-stdin
$ find /documents | feaindexadd \
  -P /ferris-index/metadata \
  --filelist-stdin
```

Listing 7.

#### A Combined Full-Text and Metadata Index Query

```
# Federation query
$ feaindexquery \
  '(&(size<=250k)(ferris-ftx==alice wonderland))'

# Recently modified local files with a given URL
$ feaindexquery \
  -P ~/clucene-index \
  '(&(mtime>=begin last week)(url=~journal))'
```

shared database, you have to add files to the full-text index first. This limitation is due to the CLucene API.

### Query Time

We now have the choice of looking in our personal files, the file server or both with our queries. The query syntax is identical for all three; we need to specify only which index to use. If we don't specify an index, we use the default, which on our desktop machine is our federation. Shown in Listing 7 are a few example queries. The =~ operator in the last example is a regular-expression match.

### Search Interfaces

libferris can present the result of a query as a filesystem. This can provide a quick interface for clients on the network to query the file server. The ferrisls command can output its results as an XML file. Given a Web form and your favourite Web scripting language, queries can be run with ferrisls, and the resulting XML file XSL translated into nice HTML for the client.

The FUSE module also allows access to search results directly through the kernel ready for exporting to the network.

The eqq:// virtual filesystem takes a query as a directory name and will populate the virtual directory with files matching the query. Other closely related query filesystems are the eaquery:// tree. The eaquery:// filesystem has slightly longer URLs, but it allows you to set limits on the number of results returned and to set how conflicting filenames are resolved. Some example queries are shown in Listing 8. Normally, a file's URL is used as its filename for eaquery:// filesystems. The short-

Listing 8.

#### Query Results as a Filesystem

```
# All files modified recently
$ ferrisls -lh "eqq://(mtime>=begin last week)"

# Same as above but limited to 100 results
# as an XML file
$ ferrisls --xml \
  "eaquery://filter-100/(mtime>=begin last week)"

# limit of 10,
# resolve conflicts with version numbers
# include the desired metadata in the XML result
$ ferrisls --xml \
  --show-ea=mtime-display,url,size-human-readable \
  "eaquery://filter-shortnames-10/(mtime>=blast week)"
```

Listing 9.

#### Alter the URLs Returned by the File Server for Local NFS Mountpoints

```
$ feaindex-federation-add-url-substitution-regex-for-index \
  --sub-index-path /ferris-index/metadata \
  --regex '^file:[/]+tmp/(.*)' \
  --format 'file:///mytmp/\1'

$ feaindexquery '(ferris-ftx==alice)'
file:///mytmp/alice13a.txt
```



names option uses only the file's name, and when two results from different directories happen to have the exact same filename, it appends a unique number to one of the result's filenames. This is likely to happen for common file names, such as README.

### When URLs Are Not Universal

The default federation plugin assumes that for any file the same URL is used to access it from all indexes in the federation. For example, consider a file with URL file://doc/lj.txt on the file server. If this file is returned as a match to a federated query, the person performing the search will want to find the file at file://doc/lj.txt relative to his or her local machine. If the /doc directory is exported as an NFS share for desktop machines, it should be mounted as /doc on the clients.

If paths between the file server and clients differ, URL modification can be done by the federation plugin. The supported URL modification will be familiar to Perl users. For each index in the federation, a regex and format string can be provided to rewrite URLs returned from that index. URL rewriting is shown in Listing 9. This example will alter any files from /tmp on the file server to be mytmp on the desktop machine.

### Caveats

In order to determine if a document has not changed since it was indexed, the PostgreSQL index plugins load some information from

the database into a RAM cache. If more than one process is updating a PostgreSQL index, more work may be done than is strictly necessary. The PostgreSQL index plugins are safe to be updating the index while clients are performing queries. Many of the other plugins provide only the level of concurrent access that the underlying index library offers. This usually amounts to many index readers or one exclusive writer.

There are Xapian index plugins for both metadata and full-text indexes. Unfortunately, Xapian has limited support for metadata queries, mainly equality only. For a metadata and full-text combination, using Xapian for both, files must be added to the metadata index first and then the full-text index.

The CLucene plugins are much easier to use than the Lucene ones. The latter relies on GCJ and an install of Lucene that GCJ can compile C++ code against.

Additional effort is required to use the PostgreSQL index plugin for a file server index that supports emblem and geospatial queries. ■

**Resources for this article:** [www.linuxjournal.com/article/9390](http://www.linuxjournal.com/article/9390).

Ben Martin has been working on filesystems for more than ten years. He is currently working toward a PhD combining Semantic Filesystems with Formal Concept Analysis to improve human-file-system interaction.


Statement of Ownership, Management, and Circulation		
1. Publication Title: <i>Linux Journal</i>	PO Box 55549	10. Owner(s):
2. Publication Number: 1075-3503	Seattle, WA 98155-0549	Phil Hughes
3. Filing Date: October 11, 2006	9. Full Names and Complete	PO Box 55549
4. Issue Frequency: Monthly	Addresses of Publisher, Editor,	Seattle, WA 98155-0549
5. Number of Issues Published	and Managing Editor:	Joyce Searls
Annually: 12	Publisher: Phil Hughes	PO Box 55549
6. Annual Subscription Price: \$25	PO Box 55549	Seattle, WA 98155-0549
7. Complete Mailing Address of	Seattle, WA 98155-0549	Adele Soffa
Known Office	Editor: Nick Petreley	PO Box 55549
of Publication:	PO Box 55549	Seattle, WA 98155-0549
1752 NW Market St. #200	Seattle, WA 98155-0549	11. Known Bondholders,
Seattle, WA 98107	Managing Editor:	Mortgagees, and Other Security
Contact Person: Jill Franklin	Jill Franklin	Holders Owning or Holding 1
206-782-7733 x112	PO Box 55549	Percent or More of Total
8. Complete Mailing Address of	Seattle, WA 98155-0549	Amount of Bonds, Mortgages,
Headquarters of General		or Other Securities: None
Business Office of Publisher:		12. Tax Status: Has not Changed
		During Preceding 12 Months
13. Publication Title: <i>Linux Journal</i>	14. Issue Date: October	
15. Extent and Nature of Circulation	Average No. Copies Each Issue	No. Copies of Single Issue
	During Preceding 12 Months	Published Nearest to Filing Date
a. Total Number of Copies:		
(Net press run)	68,254	63,753
b. Paid and/or Requested Circulation		
(1) Paid/Requested Outside-County		
Mail Subscriptions on Form 3541.	23,970	21,466
(2) Paid In-County Subscriptions		
Stated on Form 3541	0	0
(3) Sales Through Dealers and		
Carriers, Street Vendors, Counter Sales,	29,468	28,685
and Other Non-USPS Paid Distribution		
c. Total Paid and/or Requested Circulation	53,438	50,151
d. Free Distribution Outside the Mail		
(1) Outside-County as Stated		
on Form 3541	1,454	2,340
(2) In-County as Stated		
on Form 3541	0	0
(3) Other Classes Mailed Through		
the USPS	0	0
e. Free Distribution Outside the Mail	2,925	850
f. Total Free Distribution	4,379	3,190
g. Total Distribution	57,817	53,341
h. Copies Not Distributed	10,437	10,412
i. Total	68,254	63,753
j. Percent Paid and/or Requested Circulation	92%	94%

PS Form 3526

# FREE STUFF YOU ASK?

## OK!


Send us a postage-paid, self-addressed envelope to the below address and we'll return a handful of Linux and Linux Journal stickers to you free of charge.



**LINUX JOURNAL**  
 Attn: Sticker promo  
 PO BOX 980985  
 Houston, TX 77098

## Plug This in Your Pipe and Smoke It

Configurability and extensibility are defining attributes in software appeal.



Nick Petreley, Editor in Chief

If you gestalt the Editors' Choice Awards, you should discover an interesting pattern. There were a number of software winners that beat out the competition specifically because you can extend the application to suit your personal tastes.

Firefox had this category sewn up, even if other editors hadn't expressed their desire to see Firefox get Editors' Choice. But those who defined their reasons for choosing Firefox emphasized how easy it is to extend and customize it.

Thunderbird was my first choice for the win, but I might not have gone with Thunderbird if it wasn't for the fact that the only other votes I received from other editors regarding e-mail clients were in favor of Thunderbird. This time, nobody specifically stated that they liked Thunderbird because of the extensions, but I suspect that was a factor. It is certainly the main reason why I use Thunderbird when I'm not using Mutt.

Take Eclipse as another example. There is a huge repository of plugins for the Eclipse integrated development environment (IDE). You can customize it to be a great Java

development platform, C++ development platform, PHP development platform or whatever else you want it to be. This is undoubtedly why it has become the favorite IDE among professional Linux developers, according to Evans Data Corporation.

KDevelop is yet another shining example, even though it lost out to Eclipse in the category of development tools. It not only lets me customize it for different languages, but also it automatically places multiple startup configurations in the KDE menu. I can start KDevelop as a C/C++ IDE, Ruby IDE or multilanguage IDE. I don't know if this is unique to Ubuntu/Kubuntu or if this is how KDevelop installs on other distributions, but I like it.

Then there's AbiWord. AbiWord wouldn't have had a chance against the competition for Editors' Choice if it wasn't for the fact that there are so many good plugins available.

Forgetting Editors' Choice for a moment, Jedit is my personal favorite editor because I can add the features and usability enhancements that I find most appealing. I use about a dozen of the many plugins available for Jedit to customize the editor to suit exactly my tastes and needs.

I also use the Google custom home page. I like the huge assortment of gadgets and feeds from which to choose, and how you can drag them around the page and drop them where you want them. Yahoo has the same kind of customizable page. Those of us who like this sort of thing can't be in the minority. Microsoft figures that the fact that it failed to offer a customizable page is one reason why it can't catch up to Yahoo or Google in the search game. So Microsoft created a customizable home page for [www.live.com](http://www.live.com) with a twist: tabbed pages. Google responded by adding tabs as a feature for its custom home pages.

My apologies to anyone who is tired of the GNOME vs. KDE debate, but I believe this is where GNOME went wrong. Before you mail that flame, let me bring you up to date

a little. I've been giving GNOME yet another chance, and this time I actually like it. I like it quite a bit, especially the way Ubuntu pre-configures GNOME. If I end up going back to KDE, it won't be a matter of fleeing back to KDE as in the past. It will finally be a simple matter of preference, not the conclusion that GNOME is so broken as to make it unusable.

However, one of the things I still don't like about GNOME is how difficult it is to customize its look and behavior to my heart's content. Given the above, it's clear most people don't want things to just work, they want things to just work *their way*. GNOME isn't totally inflexible by any means. It lets me customize a lot, to its credit. But there are places where it is needlessly restrictive. I've been able to find ways to work around GNOME's design to get it to work my way (for the most part, anyway), thanks to some guidance from readers. But when you have to use the gconf-editor, edit files or install non-GNOME utilities to get what you want, it is obvious that this level of customization runs contrary to the design philosophy of GNOME. GNOME developers need to take a cue from Microsoft and realize how important users consider personalization and that users want the process to be as easy as possible.

People's preferences can differ greatly. One editor who gushed about how easy it is to extend Firefox to suit one's own needs named as his favorites a number of extensions that I didn't know exist. I don't want the extensions he uses, and I'm glad they don't take up space on my workstation.

I predicted long ago that software would end up this way, although at the time I thought it would take the form of networked components instead of downloadable plugins. Nevertheless, it's the right way to go, and I'm glad to see that open-source projects blazed this trail. ■

---

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](http://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](mailto: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](http://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](http://microway.com)