

# LINUX JOURNAL™

Since 1994: The Original Magazine of the Linux Community

JUNE 2006 | www.linuxjournal.com

DUAL BOOTING  
WITH FINESSE

TRIP UP INTRUDERS  
WITH TRIPWIRE

Easy, Safe  
Access to

# STORAGE

And Not So  
Easy Ways To  
Recover RAID



AN  
**SSC**  
PUBLICATION

USA \$5.00  
CAN \$6.50



RAID AND LVM2 DATA RECOVERY • SSH IS NOT JUST A SECURE SHELL  
FISH FOR DATA WITH KIO • BOOT LINUX WITH AN IPOD

# Let's go Xtreme

## Introducing Appro XtremeServers & Workstation with 8 DIMM Sockets per CPU

- 2-way or 4-way, Single or Dual-Core AMD Opteron™ processors
- **Largest memory capacity - 8 DIMM Sockets per CPU, up to 128GB**
- PCI-Express technology to increase I/O bandwidth and reduce system latency
- Outstanding Remote Management – IPMI 2.0 compliant
- Cable-free design, ready to run, simple to install, service and maintain
- Support for Windows® or Linux OS
- Ideal for memory-intensive and I/O-intensive applications



1U / 2U / 3U Servers  
and Workstation

---

AMD Opteron™ Processors - AMD64 dual-core technology reduces memory latency and increases data throughput  
- Dual-core processors with Direct Connect Architecture deliver the best performance per watt with little or no increase in power consumption or heat dissipation.

---



Appro delivers high-performance computing solutions to help you maximize productivity for a solid ROI. On-site maintenance and installation services are also available.

For more information, please visit [www.appro.com](http://www.appro.com)  
or call Appro Sales at 800.927.5464, 408.941.8100.



*The Power of Being There.<sup>®</sup>*

*Times two.*

Discover how the combined power of Avocent and Cyclades IT infrastructure management solutions can take you and your data center to the next level. KVM, serial and power – all over IP. Plus, Intelligent Platform Management Interface (IPMI) and embedded KVM. The Power of Being There.<sup>®</sup> **Times two.**



**Avocent.<sup>®</sup>**



**cyclades**

An Avocent<sup>®</sup> Company




Visit [www.avocent.com/powerx2](http://www.avocent.com/powerx2)

Avocent, the Avocent logo, The Power of Being There and Cyclades are registered trademarks of Avocent Corporation or its affiliates. All other trademarks or company names are trademarks or registered trademarks of their respective companies. Copyright © 2006 Avocent Corporation.

# CONTENTS

JUNE 2006

Issue 146



storage

## FEATURES

### 52 RECOVERY OF RAID AND LVM2 VOLUMES

When there's something strange in your LVM, who you gonna call?

Richard Bullington-McGuire

### 58 NETWORK TRANSPARENCY WITH KIO

Konqueror is a slave to fishing. Or at least it has one.

Jes Hall

### 62 YELLOW DOG LINUX INSTALLS NEATLY ON AN IPOD

Take one Mac, insert iPod, boot Linux.

Dave Taylor

### 64 SSHFS: SUPER EASY FILE ACCESS OVER SSH

SSH does more than just provide safe communications.

Matthew E. Hoskins

#### ON THE COVER

- Dual Booting with Finesse, p. 89
- RAID and LVM2 Data Recovery, p. 52
- SSH Is Not Just a Secure Shell, p. 64
- Fish for Data with KIO, p. 58
- Boot Linux with an iPod, p. 62

COVER PHOTO: JOHN LAMB/STONE+/GETTY IMAGES

# 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.



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



# CONTENTS

## JUNE 2006

### Issue 146

## COLUMNS

### 24 REUVEN M. LERNER'S AT THE FORGE

Google Maps

### 28 MARCEL GAGNÈ'S COOKING WITH LINUX


If Only You Could Restore Wine

### 34 DAVE TAYLOR'S WORK THE SHELL

Coping with Aces

### 38 MICK BAUER'S PARANOID PENGUIN

Security Features in Red Hat Enterprise 4



### 44 JON "MADDOG" HALL'S BEACHHEAD

Shoring Up the Seawall

### 46 DOC SEARLS' LINUX FOR SUITS

Use and Usefulness

### 96 NICHOLAS PETRELEY'S ETC/RANT

SUSE Rocks, Fedora Locks

## IN EVERY ISSUE

10 LETTERS

14 UPFRONT

50 NEW PRODUCTS

81 ADVERTISER INDEX

## INDEPTH

### 68 AN INTRODUCTION TO GAMBAS

Will VB refugees gamble on Gambas?

Mark Alexander Bain



### 76 HOW TO SET UP AND USE TRIPWIRE

Don't let intruders go unnoticed.

Marco Fioretti

### 80 THE WORLD IS A LIBFERRIS FILESYSTEM

libferris can make your toaster look like a filesystem.

Ben Martin

### 86 PENDRIVES AND THE DISTRIBUTIONS FOR THEM

A look at the distros you can use for booting Linux from a pendrive.

Juan Marcelo Rodriguez

### 89 THE ULTIMATE LINUX/WINDOWS SYSTEM


Some of your Windows and Linux applications can share configuration data.

Kevin Farnham

### 93 CONVERTING VIDEO FORMATS WITH FFMPEG

FFmpeg is a mini Swiss Army knife of format conversion tools.

Suramya Tomar



### 89 THE ULTIMATE LINUX/WINDOWS SYSTEM

## Next Month

## RUBY

Next month we get you started programming in a finely cut and polished language called Ruby. We also have some valuable content for those who are already Ruby programmers, like how to manage your Ruby libraries. And, did you know that you can use Ruby to glue together a vast variety of resources in your enterprise? We also explore Ruby on Rails and the changes in version 1.1. Marcel takes a whole 'nother slant on Ruby and introduces you to a variety of games with rubies in them.

There's more. Jon "maddog" Hall charts the depths of the sea, littered with the abandoned sunken hulls of proprietary design. Doc Searls addresses a related issue, runaway patents and absurd copyrights. All in the next issue of *Linux Journal*.

# TYAN



## Power of 4 + 4

### TYAN Matchless 8-Processor Server



#### M4881

- Processor Expansion board for S4881**
- Up to four AMD Opteron™ processor MP
  - Up to 64GB of Registered (ECC) DDR400/333
  - 166 DIMM sockets
  - 12"x13" dimension with special architecture
  - Two Hyper Transport™ v1.0 Connectors



#### Thunder S4881

- 8-Processor AMD Opteron™ Motherboard**
- Up to eight or four AMD Opteron™ processor
  - Up to 64GB of Registered (ECC) DDR400/333
  - (4)SATA support RAID 0, 1, 0+1
  - Broadcom® BCM5704C Dual-Channel GbE controller
  - IEEE 1394a (FireWire) ports
  - ATI™ RAGE™ XL with 8MB



#### Transport VX50 B4881

**8-Processor AMD Opteron™ 5U Server**

- 4/8 Processor HPC Computing Platform
- Support up to eight (8) AMD Opteron™ 6xx series
- 4-P model (B4001-4P):
  - Sixteen (16) DIMMs, supporting max. 64GB registered, ECC DDR400/333 memory
- 8-P model (B4881-8P):
  - Thirty-two (32) DIMMs, supporting max. 128GB registered, ECC DDR400/333 memory
- Two (2) gigabit Ethernet LANs, Broadcom® BCM5704 controller
- Three (3) 64-bit, 133/100/66MHz PCI-X slots , Two (2) PCI-E slots (1\*x16 + 1\*x4 signal)
- Total 5 usable PCI expansion slots

**TYAN COMPUTER CORP.**

#### Tyan Computer USA


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

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

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

# Zervex

Hosting Control Panel



ServerCP is the best solution for hosting companies and organizations who need a webserver, mailserver, dns server, or a fully featured web hosting environment.

## Features

1. RPM based update deployment
2. Automated Software Installer
3. Apache, PHP, MySQL
4. Multi-platform Integration

# Zervex

Announces

## Site Licensing

### \$69 / month

Check it out  
at

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

# LINUX JOURNAL

## Editor in Chief

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

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

## Contributing Editors

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

**Proofreader** Geri Gale

**VP of Sales and Marketing** Carlie Fairchild  
[carlie@ssc.com](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

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

**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 77253-3587 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 77253-3587 USA

Please allow 4-6 weeks for processing address changes and orders

PRINTED IN USA

**USPS** LINUX JOURNAL (ISSN 1075-3583) is published monthly by SSC Media Corp., 3262 Westheimer Rd, Suite 502, Houston, TX 77098. Periodicals postage paid at Seattle, Washington and at additional mailing offices. Cover price is \$5 US. Subscription rate is \$25/year in the United States, \$32 in Canada and Mexico, \$62 elsewhere. POSTMASTER: Please send address changes to *Linux Journal*, PO Box 980985, Houston, TX 77253-3587. Subscriptions start with the next issue.

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

## User Management


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

## Wireless AP and Backbone

- Wireless monitoring
  - Frequency scanning with detailed report
  - Raw wireless packet sniffer
    - streaming option to Ethereal analyzer
    - option to save to a file format supported by Ethereal
  - Snooper packet inspection
    - analyzes all raw frames received for wireless parameters
    - monitors a single channel or all channels
- Nstreme wireless polling protocol
  - no decrease in speed over long distances (as seen with the 802.11 ACK packet bottleneck)
  - polling improves speed and eliminates contention for access to the wireless bandwidth
  - access point control over Nstreme clients tx data to optimize use of the wireless medium
  - RADIUS support for the access control list including bandwidth settings for wireless clients

- Full 802.11a/b/g support

The above is a brief description of a few features, for more information and a fully featured 24 hour demo go to:



**Free**

Download at [www.mikrotik.com/dude](http://www.mikrotik.com/dude)

## "The Dude" Network Monitor

- Auto network discovery and layout
- Discovers any type or brand of device
- Device, Link monitoring, and notifications
- Includes SVG icons for devices, and supports custom icons and backgrounds
- Easy installation and usage
- Allows you to draw your own maps and add custom devices
- Supports SNMP, ICMP, DNS and TCP monitoring for devices that support it
- Link usage monitoring and graphs
- Direct access to remote control tools for device management
- Supports remote Dude server and local client
- Runs in Linux Wine enviroment, MacOS Darwin, and Windows
- Best price/value ratio compared to other products (free of charge)

## RouterBOARD 112

- Super-small embedded board
- 140mm x 85mm (5.51in x 3.35in)
- Ideal for wireless CPE solutions
- RouterOS Wireless AP/Firewall/ Routing software already included
- Best price/performance/feature ratio compared to similar products
- 175MHz MIPS CPU
- 2 miniPCI (one on each side)
- 1 10/100 Ethernet MDI-X
- 64MB NAND storage
- 16MB SDRAM
- Speaker
- Serial port
- Passive PoE (also 12V PoE)
- Power input 11-48V

**\$99**



**\$95**

- MDI-X
- 4 separate 10/100 ports



## RouterBOARD 44

### For the Router Builder

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

**\$240**

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



## RouterBOARD 500 & RouterBOARD 564

### The Wireless Switchboard AP

The bigger and more powerful RB532 accompanied with a RB564 daughterboard can be a complete multi-radio tower system, with 9 ethernets and 6 miniPCI slots.

# CIARA

**GET THE SERVER THAT  
TAKES YOUR BUSINESS AS  
SERIOUSLY AS YOU DO.**

Intel® Xeon® Processors powering  
**Ciara VXPPO™** servers provide the  
quality and dependability to keep up  
with your growing business.

## **3 Good Reasons to Buy from Ciara Technology**

### **Built To Order**

Customized Servers, Blade & Storage:  
Get the technology that's right for your business and not right for your supplier. With the capability to manufacture over 2500 systems a day, Ciara is suited to accommodate any customer requirement. Our record growth enabled us in February 2003 to inaugurate an all-new, ultra-modern manufacturing plant of 576,000 sq ft. Our systems are built under the ISO 9001-2000 certification.

### **Service...Service...Service**

Incorporated in 1984, Ciara Technology is a world-class provider of computer systems including desktop, laptop, servers, storage and supercomputer clusters as well as other software and integration services. All our systems are serviced by Ciara's highly trained and certified technicians and system engineers. We are an accountable supplier - One single point of contact for all your technology needs.

### **Channel Partner Premier**

Ciara has a strong working relationship with Intel, so we have access to information and support that give us - and you - significant advantages in deploying and managing your systems and applications. The result is a more flexible solution that meets your current needs, while enabling easy expansion to accommodate emerging technologies and new business growth.

*Visit us*

*For all your computer needs*

*Give us a call*

**CIARA-TECH.COM**

**1-866-789-7225**

# TECH

## DUAL CORE RACKMOUNT SERVER



Power Efficient  
Intel® Xeon® LV  
**120Watts**  
per server

**\$1099<sup>1</sup>**

**DUAL CORE**

### **VXPRO-R7230NH** **Affordable Starter Server**

Ultra Affordable  
Easy to Service

**Intel® Pentium® D Processor 920**  
**(2.8GHz Dual Core Processor)**  
800MHz Front Side Bus  
2x2MB L2 Cache

1GB (2x512MB) ECC DDR2 533 Memory  
(Expandable to 8GB)

One 80GB (7,200RPM) SATA150 HDD  
(Up to 4 Hot Swap Hard Drives)  
Floppy and CD-ROM included

1U Rackmount  
3 Years Warranty Return to Depot

#### **Upgrade Your System**

Intel® Pentium® D Processor 930 (3.0GHz)	<b>Add \$99</b>
Intel® Pentium® D Processor 940 (3.2GHz)	<b>Add \$140</b>
Additional 80GB (7,200RPM) SATA150	<b>Add \$72</b>
Additional 1GB (2 x 512MB) ECC DDR2 533	<b>Add \$109</b>

**\$2955<sup>1</sup>**

**DUAL CORE**

### **VXPRO-R7520BB2** **Low Voltage - 4 Cores Server**

Extremely Powerful  
Ultra Low Power Consumption

**Dual Intel® Xeon® LV Processor**  
**(2.0GHz Dual Core Processor)**  
667MHz Front Side Bus  
2MB of Shared L2 Cache

2GB (2x1GB) ECC/Reg DDR2 400 Memory  
(Expandable to 16GB)

One 80GB (7,200RPM) SATA150 HDD  
(Up to 4 Hot Swap Hard Drives)  
Floppy and CD-ROM included

1U Rackmount  
3 Years Warranty Return to Depot

#### **Upgrade Your System**

Additional 80GB (7,200RPM) SATA150	<b>Add \$72</b>
Additional 2GB (2 x 1GB) ECC/Reg DDR2 400	<b>Add \$299</b>



Coloron, Coloron Inside, Centrino, Centrino Logo, Core Inside, Intel, Intel Logo, Intel Core, Intel Inside, Intel Inside Logo, Intel SpeedStep, Intel Viv, Itanium, Itanium Inside, Pentium, Pentium Inside, Xeon and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. (1) Important Information. All prices, specifications and promotional offers are subject to change without notice. Cisco cannot be responsible for typographical errors, photographic errors, pricing errors. All pricing in US dollar. Shipping and applicable taxes are not included.



## Worth the Wait

Starting with your first /etc/rant column, I have loved them. You are saying aloud what many in the community think, with a very loud voice. I wanted to ask you, why is this available only to subscribers? I'd love to make some comments on my blog and then link to your columns, or translate them to Spanish and let my friends read them in my blog. Tell me if this is feasible, please. I already commented on the /etc/rant of February 2006 issue and plan to do it on the April one I've just received, because I have lots of friends using GNOME just because they think it embodies the spirit of Free Software. Keep up the rant...er, the work.

—  
**Enrique Verdes**

*We do not publish in print and on the Web simultaneously, but we do publish our print content on the Web after a period of time expires. You can find archives of the magazine at the URL [www.linuxjournal.com/xstatic/magazine/archives](http://www.linuxjournal.com/xstatic/magazine/archives).—Ed.*

## Rant about Rant

Teams work very hard to release and offer their programming works of art. I doubt any of them wishes to bash each other's software. How about some constructive rants? The developers just might be inclined to offer a feature/fix that your "opinion" had a problem with. Or, heck, maybe you might even spawn a reader to contribute constructively instead of complain. Thanks for listening to my rant.

—  
**Ryan Ferguson**

## Content Must Be Free of Unwanted Influence

The position of editor in chief traditionally allows ultimate control of all aspects of a publication—from which letters will be included, to whether GNOME or KDE screenshots are used, to what articles are published and who is allowed to contribute. Perhaps *LJ* has a system of checks and balances in place that the casual reader doesn't know about, but from the cheap seats, I'm concerned. Not to sound alarmist, but a strong bias from the editor in chief is a contamination that can't possibly be quarantined only to the last page of a publication!

—  
**Matt McElheny**

*It is the job of an editor in chief to serve the magazine's readers, period. That means keeping the content separate from the influence of advertising and unaffected by the editor's own personal opinions. If at any point you believe the content reflects otherwise, I will take your complaints very seriously.—Ed.*

## On the Spirit of Open Source

Great rant [April 2006]. But, as I see it, the real problem is not about advocacy, it's about (never-ending) fragmentation. Big players (such as Oracle and Dell) are complaining about it, and no one seems to listen. The real problem is that a genius like Miguel is wasting his time with GTK or Mono (the C# equivalent of the GNU Java compiler), when he could do really useful stuff. It's a pity.

—  
**Dani**

## GCJ Deserves an Apology

I just got the latest *LJ* in the mail today. I wonder about the genesis of the "practically useless GCJ" remark in the etc/rant column [April 2006].

The version of Eclipse on my Debian Sid installation was built with GCJ, and it seems to run fine with no Sun JRE in sight. I thought Eclipse was a fairly large and complex project (although I have never tried to build it from scratch myself), so what does GCJ need to be able to do to get out of the "practically useless" category?

I also thought the limits of GCJ were due to limits of the GNU classpath, and that work on that was progressing nicely. Maybe I'm not following developments closely enough, though.

—  
**Jon**

*I'm the one not following the developments closely enough. Granted, it's not supported for Swing, but if you can compile the SWT-based Eclipse with GCJ, I owe it an apology.—Ed.*

## Even Einstein Agrees

So far, I just want to say that I like the rants closing the new *Linux Journal* issues, although a blog of some kind might be a more appropriate place.

Still, you are echoing sentiments that I have expressed many times in the past. From the moment Miguel said, "UNIX sucks", and started making Linux look like Windows, I've felt like I've been lost in some kind of terrible nightmare, unable to awaken. I love UNIX, and although it certainly needs help on the desktop from a development API standpoint, I'm not about to throw away everything about UNIX that makes it great. Please reference the late Mr Einstein's genius (I use the following quotation in my e-mail signature) for why UNIX is great, IMHO:

"Any intelligent fool can make things bigger and more complex....It takes a touch of genius—and a lot of courage—to move in the opposite direction."—Albert Einstein

—  
**Michael P. Soulier**



# Remember when the sky was the limit?

With Intel® Software Development Products, the Swinburne Center for Astrophysics is showing today's kids a universe filled with unlimited possibilities.



## LINUX\* APPLICATIONS HELP UNRAVEL THE ORIGINS OF THE UNIVERSE.

Swinburne University's Center for Astrophysics and Supercomputing in Melbourne, Australia, is helping develop the next generation radio telescope in order to collect enough data to perform modeling and simulations of our entire galaxy. Their goal is to make realistic, 3D, virtual-reality animations available to the general public, particularly school children. To do this, they used the Intel® compilers to deliver application performance improvements — saving them valuable development time and money<sup>1</sup>. Whether you build applications for physics or financial analysis, Intel® Software Development Products help your Linux\* applications reach for the stars.

**Download a trial version today.**  
[www.intel.com/software/products/nolimits](http://www.intel.com/software/products/nolimits)

Copyright © Intel Corporation, 2006. All rights reserved. Intel, the Intel logo, are trademarks of Intel Corporation or its subsidiaries in the United States and other countries. \*Other names and brands may be claimed as the property of others. <sup>1</sup> Performance and benchmark information were provided by Swinburne University. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing.

Do you take

"the computer  
doesn't do that"

as a personal challenge?

So do we.

**LINUX**  
**JOURNAL**<sup>™</sup>

Since 1994: The Original Monthly  
Magazine of the Linux Community

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

## [ LETTERS ]

### **LJ Is Way Too Practical**

Please cancel my subscription, because the *Linux Journal* disappoints me thoroughly. It's too practical for this budding theoretical computer scientist.

—  
**William J. McEnaney Jr.**

### **Size Does Matter**

You now have a very inconvenient size and it ruins the display I have of *Linux Journal*. It doesn't fit on my lap in the men's reading room; it doesn't fit in most places while reading, except in your hands, and it is awkward there as well. I love *Linux Journal*; I hate the new size of the pages.

—  
**Pete Gandy**

### **SSH Quibbles**

In "Demons Seeking Dæmons—A Practical Approach to Hardening Your OpenSSH Configuration" [March 2006], Phil Moses mentions the UsersAllow directive, but it is really the AllowUsers directive (as Listing 2 shows). And UsersDeny is really DenyUsers.

The meaning of an entry such as user@hostname.domain is misleading. He seems to indicate that it allows access to user from hostname.domain, but it really means that people from hostname.domain can access user's account (according to O'Reilly's SSH book).

—  
**Byron Rendar**

### **Ow, No OWFS?**

In your recent article on temperature monitoring [see "Remote Temperature Monitoring with Linux" by Steven M. Lapinskas, *LJ*, April 2006], you didn't mention Paul Alfille's OWFS Project ([owfs.sourceforge.net](http://owfs.sourceforge.net)).

OWFS already has been ported to the Linksys WRT54G wireless router, providing a cheap and readily available hardware platform for monitoring projects like the one described in the article ([owfs.sourceforge.net/WRT54G.html](http://owfs.sourceforge.net/WRT54G.html)).

—  
**Steve**

### **Xoops! A Security Hole**

Juan Marcelo Rodriguez's XOOPS article [April 2006] tells us to set three directories so that any local user, including the Web server, can create and execute programs in them:

```
chmod 777 uploads cache templates_c
```

That's poor security practice. If XOOPS needs it, XOOPS needs a redesign. These wide open directories are just the place to install a 'bot for spamming or DoS attacks. It's begging for an Internet worm, if one doesn't exist already.

Mambo has similar issues. It seems the easier they make these complex Web apps to install, the less care they pay to securing those installations.

—  
**Cameron Spitzer**

**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.

**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).

## Data Center Solutions Built on OPEN STANDARDS

### OSS-1000 1U OPEN Series

- Available with 1, 2, 4, or 8 CPU's per system
- Up to 2 terabytes storage and 64GB of RAM
- Tool-less 1U case design with rail kits
- 4 Hot-swap drive bays, SCSI/SATA-2
- Supports RAID levels 0,1, 5, 6, and 10
- 3 year next day warranty within the USA



### OSS-2000 2U OPEN Series

- Available with 1, 2, 4, or 8 CPU's per system
- Up to 6 terabytes storage and 64GB of RAM
- Tool-less 2U case design with rail kits
- 12 Hot-swap drive bays, SCSI/SATA-2
- Supports RAID levels 0, 1, 5, 6, and 10
- 3 year next day warranty within the USA

### OSS-3000 3U OPEN Series

- Available with 1, 2, 4, or 8 CPU's per system
- Up to 8 terabytes storage and 64GB of RAM
- Tool-less 3U case design with rail kits
- 16 Hot-swap drive bays, SCSI/SATA-2
- Supports RAID levels 0, 1, 5, 6, and 10
- 3 year next day warranty within the USA



**866-664-STOR**  
**Sales@OSShpc.com**  
**OpenSourceStorage.com**



All Categories    
Shop Online for OSS on [whaBAM.com](http://whaBAM.com)



;; 1195 Borregas Avenue, Sunnyvale, CA 94089 ;;

## diff -u

WHAT'S NEW  
IN KERNEL  
DEVELOPMENT

With the imminent advent of version 3 of the **GNU General Public License**, the question naturally comes up whether Linux will continue to be licensed under version 2, or migrate to version 3. The answer seems to be that no migration will occur. Even if Linus Torvalds wanted to, he does not control the copyright of all the thousands upon thousands of code contributions that have gone into the kernel over the years, and only those copyright holders could authorize a license change for their own contributions. Although attempts to track down all the contributors and obtain permission have been made by optimistic souls, it is virtually impossible to do this. The GPL version 2 will be the Linux kernel license for the foreseeable future.

Meanwhile, more and more kernel variables and functions are being set to operate only with GPLed third-party drivers. The kernel can do this by testing any driver to see if it sets a variable indicating the license under which its code is released. If the license is the GPL, the kernel allows access to those restricted symbols. Otherwise, it's denied. This recently bit **AVM**, when **Greg Kroah-Hartman** restricted the **USB subsystem** to operate only with GPLed drivers. AVM always had released its own binary-only driver for its hardware, but this new change stopped it in its tracks. The change itself was reverted out of the kernel, although it did turn out that one reason Greg had implemented that change was because it has been possible for some time now to write USB drivers in user space with no loss of speed. Regardless, one result of this particular confrontation was for Greg to implement some logging infrastructure, so that symbols soon to be GPL-restricted would be clearly identified in the logs at runtime, and sysadmins on those systems could begin to make reasonable preparations for that change.

**Willy Tarreau** has begun gathering together useful **2.4 patches** and making them available at a central location. In general, there has been an outcry among some users that without the even/odd development dichotomy, and with 2.4 virtually stationary, there is now no stable series that is still reasonably up to date with current features. Arguments by the kernel developers that distributions take care of kernel stability, and that the w.x.y.z tree was created specifically to address the issue of stability, are unconvincing for one key reason.

Although it's true that those kernels may get good uptime, their behaviors are still poten-

tially inconsistent from version to version—that is, the code base itself is unstable, making it difficult for user-space developers to create systems that behave reliably for the services they require. Although it seems clear to me that eventually kernel development will have to give stability first-class consideration, so far there is no real movement in that direction among the top developers. This may be one reason why Willy's foray back toward 2.4 maintenance has come about.

The **IDE driver** may be going away at some point in the medium future, as **libata** becomes more and more robust, and a more viable replacement for it. **Alan Cox** is confident that, although now is still not the time for a straightforward replacement, it is still the goal and the intention of ongoing libata development. It's important to bear in mind, when considering this, that the **IDE** nightmare cannot be ended simply by replacing one set of code with another. The IDE standards are still (and will continue to be) a horrible mess, and vendor interpretation and compliance with those standards is still (and will continue to be) extremely nuanced and difficult to support. So whatever the future of IDE may be, it will have to accommodate all the twists, turns, bumps and punctures endured in the past. Even if all future IDE hardware magically conformed to a single sane standard, it would be quite some time before we could abandon support for all the older hardware.

The **Reiser 4** filesystem has not unexpectedly run into problems being accepted into the Linux kernel. After the most recent flame war, which had some developers throwing up their hands and saying they refused to give feedback on ReiserFS patches until **Hans Reiser** stopped attacking them, responses to ReiserFS posts on the mailing list have thinned out. Without the support of the kernel developers, the possibility of Reiser 4 going into the kernel becomes more problematic. Only the kernel developers fully grasp the requirements that any given patch must meet in order to be accepted into the tree. Without their guidance, Reiser 4 may have a hard time moving in the right direction. And as Reiser 4 development continues to diverge from the kernel proper, the patch that must ultimately be accepted into the kernel grows as well, adding much to the work required for final integration. Almost certainly Reiser 4 will make it into the kernel eventually, but probably not before its developers solve the technical and social issues that confront them.

—Zack Brown

## On the Web

DEBUT OF NEW  
RUBY COLUMN

Given Ruby's recent surge in popularity, linuxjournal.com is excited to bring readers The Gemcutters Shop (TGS), a new column that will focus on the Ruby community and a variety of Ruby programming topics. Regular linuxjournal.com author and active Ruby contributor Pat Eyler will show us how to develop skills related to Ruby programming, use libraries from Ruby's standard library, work with additional libraries and use applications written in or for Ruby.

Don't miss Pat's kick-off column, "Welcome to the Gemcutters Shop", available at [www.linuxjournal.com/article/8921](http://www.linuxjournal.com/article/8921). Upcoming columns will use code and case studies to demonstrate rcov, Rake, RubyGems, Mr. GUID and much more.

## 2006 EVENT CALENDAR

We're in the middle of tradeshow mania right now—pick a topic and a city, and we're pretty sure some sort of tech event is on its way. Use our Linux Industry Events calendar ([www.linuxjournal.com/xstatic/community/events](http://www.linuxjournal.com/xstatic/community/events)) to stay on top of all this year's tradeshow and conferences, from ISPCON Spring 2006 to the USENIX LISA Conference.

BE A LINUXJOURNAL.COM  
CONTRIBUTOR

Do you have a great tutorial or how-to article you want to share with an eager audience? Are you looking to unveil your coding masterpiece or helpful tricks? Did you test drive the latest version of a top-three distro and want to tell us what you really think? We're always looking for interesting and unique article proposals. Send your ideas to [webeditor@ssc.com](mailto:webeditor@ssc.com).

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

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

**Supported:**  
Media cards:  
Compact Flash,  
Secure Digital

**Supported:**  
Internal optical drive:  
CDRW, DVD±RW

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

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

**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.



# CIARA

## PEACE OF MIND

With the **Dual-Core Intel® Xeon®** Processor in your **VXRACK™**, your servers have advanced reliability features built in.

### 3 Good Reasons to Buy a VXRACK™ Cluster Technology

#### Built To Order

Customizes Servers, Blade & Storage:  
Get the technology that's right for your business and not right for your supplier. With a capability to manufacture over 2500 systems a day, Ciara is suited to accommodate any customer requirement. Our record growth enabled us in February 2003 to inaugurate an all-new, ultra-modern manufacturing plant of 576,000 sq. Our systems are built under the ISO 9001-2000 certification.

#### Service...Service...Service

Incorporated in 1984, Ciara Technology is a world class provider of computer systems including desktop, laptop, servers, storage and supercomputer clusters as well as other software and integration services. All our systems are serviced by Ciara's highly trained and certified technicians and system engineers. We are an accountable supplier - One single point of contact for all your technology needs.

#### Channel Partner Premier

Ciara has a strong working relationship with Intel, so we have access to information and support that give us - and you - significant advantages in deploying and managing your systems and applications. The result is a more flexible solution that meets your current needs, while enabling easy expansion to accommodate emerging technologies and new business growth.

*Visit us*

*For all your computer needs*

*Give us a call*

**CIARA-TECH.COM**

**1-866-789-7225**

# CIARA TECH

## HIGH DENSITY BLADE SERVER

**\$3835<sup>1</sup>**

**DUAL CORE**



### **VXBLADE-7520JR2** A Standard in the Datacenter

Powerful and Reliable  
Mid to High End Server

**Dual Intel® Xeon® Processor**  
**(2.8GHz Dual Core Processor)**  
800MHz Front Side Bus  
2x2MB L2 Cache

2GB (2x1GB) ECC/Reg DDR2 533  
(Expandable to 12GB)

One 80GB (7,200RPM) SATA150 HDD

3 Years Warranty Return to Ciara

#### Upgrade Your System

Upgrade to 120GB (7,200RPM) SATA150  
Additional 2GB (2 x 1GB) ECC/Reg DDR2 400  
Upgr. to 8GB (4 x 2GB) ECC/Reg DDR2 400

**Add \$35**  
**Add \$299**  
**Add \$899**

**\$2755<sup>1</sup>**

**DUAL CORE**

Power Efficient  
Intel® Xeon® LV  
**120Watts**  
per server



### **VXBLADE-7520BB2<sup>2</sup>** Low Voltage - Quad Cores Server

Extremely Powerful  
Ultra Low Power Consumption

**Dual Intel® Xeon® LV Processor**  
**(2.0GHz Dual Core Processor)**  
667MHz Front Side Bus  
2MB of Shared L2 Cache

2GB (2x1GB) ECC/Reg DDR2 400  
(Expandable to 16GB)

One 80GB (7,200RPM) SATA150 HDD

3 Years Warranty Return to Ciara

#### Upgrade Your System

Upgrade to 120GB (7,200RPM) SATA150  
Additional 2GB (2 x 1GB) ECC/Reg DDR2 400  
Upgr. to 8GB (4 x 2GB) ECC/Reg DDR2 400

**Add \$35**  
**Add \$299**  
**Add \$899**



### **VXR-72 or VXR-96** Special Rack Enclosure

VXR-96 accomodates the following:  
Up to 8 enclosure chassis  
Up to 96 VXBLADE  
1.1TB of memory  
48 TB of local storage

VXR-72 accomodates the following:  
Up to 6 enclosure chassis  
Up to 72 VXBLADE  
0.8TB of memory  
36 TB of local storage

Optional Vertical Cable Management  
Optional Infiniband Cable Management  
Optional Color customization



# LJ Index, June 2006

1. Number of minutes per day Britons spend watching TV: **148**
2. Number of minutes per day Britons spend using the Net: **164**
3. Millions of Weblogs as of March 9, 2006: **30.5**
4. Millions of RSS feeds, worldwide: **70**
5. Number of generations back that all humans have the same ancestors: **120**
6. Trillions of dollars in the latest US federal budget: **2.8**
7. US federal credit limit, in trillions of dollars: **9**
8. Total billions of searches by Americans in January 2006: **5.48**
9. Millions of households promised fiber broadband by 2006: **86**
10. Promised symmetrical fiber performance in Mbps by 2006: **45**
11. Average monthly dollar price promised for home fiber service by 2006: **50**
12. Dollar prices for 100Mbps service found in Japan and Korea today: **40**
13. Worldwide position of US in broadband deployment in 1996: **1**
14. Worldwide position of US in broadband deployment in 2005: **16**
15. Estimated billions of dollars lost to customers of failed US fiber deployments: **200**
16. Estimated trillions of dollars lost to the US economy by failed fiber deployments: **5**
17. Percentage of local authorities using Linux in the UK: **33**
18. Percentage of local authorities using Linux in France: **71**
19. Percentage of local authorities using Linux in Holland: **55**
20. Percentage of local authorities using Linux in Germany: **68**

**Sources:** 1, 2: *Guardian Unlimited*, from a Google survey | 3: Technorati | 4: SocialText | 5: *Slate and Nature* | 6, 7: *Washington Post* | 8: Center for Media Research | 9–16: CAZITech, TeleTruth.org | 17–20: John Dwyer in *Industry* (UK)  
—Doc Searls

## They Said It

**Don't build your own kernel. It's a great way to waste a month of your time.**

The main reason for using BSD is that you work at Yahoo. Otherwise, use Linux.

Premature optimization is the root of all evil.

—Cal Henderson, on building Flickr on Linux. Flickr is now owned by Yahoo, most of which runs on Yahoo's BSD ([conferences.oreillynet.com/cs/et2006/view/e\\_sess/8068](http://conferences.oreillynet.com/cs/et2006/view/e_sess/8068))

**REST is the Unix pipe of the net.**

—Kevin Marks, on IRC at a conference

**Daddy, you have a picture of the Internet on your shirt.**

—Six-year-old daughter of Phil Windley, on his lap while he wore a Firefox shirt

**Open source development violates almost all known management theories.**

—Marletta Baba, Dean of the College of Social Sciences, Michigan State University (source: Greg Kroah-Hartman in [www.kroah.com/linux/talks/oscon\\_2005\\_state\\_of\\_the\\_kernel](http://www.kroah.com/linux/talks/oscon_2005_state_of_the_kernel))

**Keep Your Exits Open: How Startups & Their Investors Can Minimize the Risk of Using Open Source Code**

—Title of a Dow Jones Virtual Seminar on March 20, 2006

**Make Business Fast, Easy and Risk-Free: What Free and Open Source Software Does to Liberate Free Markets**

—Title of a virtual seminar proposed by Doc Searls to Dow Jones (Source: Andrew Leyden and Doc Searls)

# Thinkmate server and workstation products offer more configurations, and more customizable options than any other system builder.



Thinkmate 1U Server, 2x Serial-ATA or SCSI drives with Intel® Xeon® Processors

## Dozens of customizable systems online, unlimited possibilities by phone.

- ▶ Rackmount Servers
- ▶ High-Performance Servers
- ▶ Storage Servers
- ▶ Pedestal Servers
- ▶ Silent Workstations
- ▶ Blade Servers
- ▶ Clustering
- ▶ Notebooks

### Operating System Options:

Thinkmate systems are available with either No OS, or pre-loaded with Linux or Microsoft Windows operating systems. Thinkmate also offers dual-boot and virtualization options.

### 100% True Hardware Customization:

Thinkmate is an innovative provider of an extensive variety of computer solutions. We completely customize all of our machines to match your individual needs. Our online quoting and ordering system has more customizable options than any other system builder on the web. If you can't find exactly what you need on our site, then give us a call and we would be more than happy to help you find it!

### Service:

Thinkmate takes customer service to a new level. All of our systems have a minimum of a 3-year advanced replacement warranty and offer up to a 3-year next business day onsite warranty through IBM Global Services. We understand mission critical situations and provide superior services to keep all of our customers satisfied.

### GSA Scheduling:

We offer rapid GSA scheduling for custom configurations. If you have a specific hardware requirement, we can have your configuration posted on the GSA schedule within 2-4 weeks.



# THINKMATE

**(800) 371-1212**  
**WWW.THINKMATE.COM**



# Monarch Computer.com™



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



**nero**  
7 ULTRA EDITION

with an upgrade to AMD Athlon™ 64!

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

**NEW!** Enter to win 1 of 8 Autographed Lance Armstrong Jerseys!



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

**FREE INSTALLATION SETUP & TESTING BY CERTIFIED TECHS**

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

## AMD Motherboard Combos

DFI NF4 Infinity nF4 w/ AMD Opteron™ processor 146 (939 pin)  
**Only \$249**

Abit AN8-32X nF4 w/ AMD Opteron™ processor 185 (939 dual-core)  
**Only \$1325**



Gigabyte GA-K8NF-9 nF4 Audio/GB-Lan/PCI-E with AMD Athlon™ 64 processor 3500+ (939)  
**Only \$275**

Asus A8R32-MVP Deluxe CrossFire w/ AMD Athlon™ 64 processor FX-60 (939 dual-core)  
**Only \$1209**


Tyan S2882G3NR-D w/ 2 x AMD Opteron™ processor 246 (940 pin)  
**Only \$719**

Tyan S2895UA2NRF K8WE w/ 2 x AMD Opteron™ processor 265 (940 dual-core)  
**Only \$1169**

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

## AMD Barebone Systems

Go to [www.monarchcomputer.com](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 Case w/450W PS  
DFI NF4 Infinity nForce4 MB  
AMD Opteron™ processor 146 (940)1.6 GHz  
**Starting @ \$395**

SuperMicro (Barebone) 1020A-T w/Hot-Swap SATA & 500W PS  
2 x AMD Opteron™ processor 246 (940 pin) & 2 GB Corsair REG ECC DDR  
**Starting @ \$1449**

Cooler Master Praetorian 730 w/535W SLI PS  
Abit AN8-32X nForce4 Motherboard  
AMD Opteron™ processor 165 (939 Dual-Core)  
**Starting @ \$755**

Cooler Master RC-810 CM Stacker 810 w/460W PS  
Tyan S2882G3NR-D Motherboard  
2 x AMD Opteron™ processor 250 (940)  
**Starting @ \$1309**



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

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

## AMD Price Drop!

**AMD Athlon™ 64 X2 Retail CPUs**  
AMD Athlon™ 64 X2 3800+ \$297.00  
AMD Athlon™ 64 X2 4200+ \$357.00  
AMD Athlon™ 64 X2 4400+ \$460.00  
AMD Athlon™ 64 X2 4600+ \$547.00  
AMD Athlon™ 64 X2 4800+ \$632.00

**AMD Athlon™ 64 Retail Box CPUs**  
AMD Athlon™ 64 3200+ (939) \$154.00  
AMD Athlon™ 64 3500+ (939) \$203.00  
AMD Athlon™ 64 3700+ (939) \$235.00  
AMD Athlon™ 64 3800+ (939) \$284.00  
AMD Athlon™ 64 4000+ (939) \$336.00  
AMD Athlon™ 64 FX-57 (939) \$811.00  
AMD Athlon™ 64 FX-60 (939) \$1011.00

**AMD Opteron™ Dual Core OEM CPUs (939 pin)**  
AMD Opteron™ 165 1.8GHz \$331.00  
AMD Opteron™ 170 2.0GHz \$402.00  
AMD Opteron™ 175 2.2GHz \$514.00  
AMD Opteron™ 180 2.4GHz \$707.00  
AMD Opteron™ 185 2.6GHz \$1134.00

**AMD Opteron™ OEM CPUs (939)**  
AMD Opteron™ 144 1.8GHz (939) \$171.00  
AMD Opteron™ 146 2.0GHz (939) \$213.00  
AMD Opteron™ 148 2.2GHz (939) \$261.00  
AMD Opteron™ 150 2.4GHz (939) \$331.00  
AMD Opteron™ 152 2.6GHz (939) \$716.00  
AMD Opteron™ 154 2.8GHz (939) \$910.00

**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™ 248 2.2GHz \$677.00  
AMD Opteron™ 850 2.4GHz \$847.00  
AMD Opteron™ 852 2.6GHz \$1130.00  
AMD Opteron™ 854 2.8GHz \$1469.00

**AMD Opteron™ Dual Core OEM CPUs (Low-wattage) OEM CPUs**  
AMD Opteron™ 265 1.8GHz \$307.00  
AMD Opteron™ 270 2.0GHz \$441.00  
AMD Opteron™ 275 2.2GHz \$689.00  
AMD Opteron™ 280 2.4GHz \$825.00  
AMD Opteron™ 285 2.6GHz \$1019.00  
AMD Opteron™ 865 1.8GHz \$677.00  
AMD Opteron™ 870 2.0GHz \$847.00  
AMD Opteron™ 875 2.2GHz \$1130.00  
AMD Opteron™ 880 2.4GHz \$1469.00  
AMD Opteron™ 885 2.6GHz \$2085.00

**AMD Opteron™ Dual Core HE (Low-wattage) OEM CPUs**  
AMD Opteron™ 260 HE 1.6GHz \$441.00  
AMD Opteron™ 265 HE 1.8GHz \$689.00  
AMD Opteron™ 275 HE 2.0GHz \$1019.00  
AMD Opteron™ 860 HE 1.6GHz \$847.00  
AMD Opteron™ 865 HE 1.8GHz \$1130.00  
AMD Opteron™ 870 HE 2.0GHz \$1469.00

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

## Components and Upgrades 1000s of In-Stock Components

101051 Cooler Master RC-830-SSN1 CM Stacker Aluminum Case w/ No FS  
**\$279.00**

150353 Western Digital Raptor X 150 GB 16 MB Cache SATA 10K (WD1500AHFD) - Clear Cover  
**\$336.00**

100828 ATI (HIS) Radeon X1800 GTO 256MB DDR3/PCI-E/AVIVO Dual-DVI-DL (Retail Box)  
**\$235.69**

140310 1 GB (2 pcs 512) DDR (400) PC-3200 OCZ Copper (OCZ4001024PF0C-K)  
**\$116.61**

101062 Antec NeoHE 500 ATX 2.01 500W Power Supply  
**\$98.77**

150939 Western Digital Caviar SE 250 GB SATA3G 16MB Cache 7200 RPM (WD2500K3)  
**\$96.90**

150269 Adaptec 2169900R 2820SA 8-Channel SATA3G PCI-X RAID Controller (Retail Box)  
**\$465.85**

140803 1 GB (2 pcs 512) DDR (400) PC-3200 Corsair XMS (TwinX1024-3200C2PT)  
**\$84.55**



**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 Linux. After properly installing the appropriate operating system, users must enable the protection of their applications and associated files from buffer overrun attacks. Consult your OS documentation for information on enabling EVP. Contact your application software vendor for information regarding use of the application in conjunction with EVP. AMD and its partners strongly recommend that users continue to use a third party anti-virus software as part of their security strategy. Please see [monarchcomputer.com](http://monarchcomputer.com) for important contract terms and conditions. No purchase required for entry.

**Sec. 508**

**NEW NEW NEW:**  
**Monarch Systems now Available with Section 508 Documentation (VPATs)!**



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

**DON'T JUST TAKE OUR WORD FOR IT!**  
**Compare our prices to these other manufacturers:**

Manufacturer & Model	Processor	Memory	Hard Drive	Their Price:	Monarch Empro:	You Save:
HP ProLiant DL145 G2	Opteron 246 x 2	2 GB	80 GB SATA	\$2,556.00	<b>\$1,778.89</b>	<b>\$777.11</b>
Sun Fire V20z Server	Opteron 248 x 2	2 GB	73 GB 10k SCSI	\$3,295.00	<b>\$2,496.27</b>	<b>\$798.73</b>
HP ProLiant DL 145 G2	Opteron 275 x 2	2 GB	36 GB 15k SCSI	\$4,766.00	<b>\$3,881.87</b>	<b>\$884.13</b>
IBM eServer 326m	Opteron 252 x 1	1 GB	80 GB SATA	\$1,989.00	<b>\$1,780.90</b>	<b>\$208.10</b>
Alienware MJ-12 7750a	Opteron 252 x 2	2 GB	250 GB SATA	\$5,456.00	<b>\$4,930.89</b>	<b>\$525.11</b>

**Monarch Saves You an Average of \$638.64!**

Pricing and data for comparison was compiled at each manufacturer's ecommerce website on December 5, 2005. All trademarks are property of their respective owners. Monarch Computer Systems, Inc. is unaffiliated with and does not make or license any products of the competing manufacturers that appear in this comparison.

**Monarch 1U Empro® 4-Way Custom Rack Server (PCI-E)**



ask for Part 80675:  
 Case/Motherboard - SuperMicro (Barebone) AS1040C-8 1U 4-Way Dual-Core Opteron W/Hotswap SCSI & 1000W PS  
 From 1 to 4 AMD Opteron™ CPUs - single or dual-core, depending on configuration  
 Extensive memory options  
 Select hard drives not only by size or interface, but by the top name brands as well - NO surprises - YOU choose what hardware powers your company.  
 Multiple RAID Setups available  
 Industry Standard Upgradable  
 1 or 3 year warranties available  
 24/7 on-site service available  
 Net Terms available  
 Dedicated Sales and Support available

**Features 4-way Dual-Core AMD Opteron™ processors**

**Sarts @ Only \$1695!**



**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.



# Saving Dollars with CentOS

At the Austin BarCamp in March 2006, I found myself sitting next to a guy with a large new laptop that seemed to be running some kind of Linux. Turned out his name was Matt Lawrence, his laptop was a new Dell Inspiron 9300, his job was systems administration and his Linux of choice was CentOS.

I asked him about it. Here are my notes, all quotes:

CentOS's goal is to be as close to Red Hat as possible without violating trademarks and copyrights. They lag behind Red Hat Enterprise 4 by two or three weeks. All the security and other updates are constantly coming out.

If you want support from a person, buy Red Hat Enterprise Linux. If you want support from a community with the same code base, get CentOS.

They fit nicely together. You can run Red Hat Enterprise for production and run CentOS for development and testing.

It's my preferred desktop. And I'd say it's an excellent desktop system choice for



Matt Lawrence with his Dell Running CentOS

small companies. Getting it up and running on this new Dell was easy. It came right up, ready to go. The only glitch was X. I had to edit the X config file for screen, which is 1920 x 1200. That's 16 x 9 rather than 4 x 3.

So I got the minimum up off one CD, then added the packages I wanted. It took about an hour to pull those down over DSL using yum.

It's a good deal. I get paid for doing Red Hat. And I save my own money

running CentOS.

We'd appreciate hearing about your experiences with CentOS as well.


—Doc Searls

## A New J for the LAMP Stack

From Tom Limoncelli (an engineer with Google and co-author of *Time Management for System Administrators*) comes an enthusiastic welcome for Jifty, a Web application development framework. Tom says Jifty is "like a Ruby on Rails for Perl". It comes from from Jesse Vincent, David Glasser and Alex Vandiver. Jesse will be familiar to fans of Request Tracker—an open-source CRM. Find it at [Jifty.org](http://Jifty.org) and [CPAN.org](http://CPAN.org).

—Doc Searls

USER FRIENDLY by J.D. "Illind" Frazer



LINUX JOURNAL EDITION



# iSCSI that makes your IP SAN a whole lot richer...





## SBE's feature-rich iSCSI software solutions deliver enterprise-level reliability at affordable price points...

### Highly scalable & fully standards compliant iSCSI Target and Initiator protocol stacks.

Early iSCSI deployments did not offer the high-quality transport and fault tolerance needed by enterprise storage managers. However, SBE's PyX iSCSI solutions were among the industry's first to provide the full error recovery features required by enterprise storage managers, namely Error Recovery Level 2 (ERL2). Providing broad Linux OS support, SBE's iSCSI protocol stack enables robust, cost-effective IP-based storage transport solutions delivering fault-tolerant, high availability networks with a rich mix of advanced feature sets, including:

- ▶▶ Proven performance from Wi-Fi to 10G with unlimited port count
- ▶▶ Aggregation of all available bandwidth across subnets and ports
- ▶▶ Multi-pathing I/O functionality for maximum redundancy and reliability
- ▶▶ Performance as fast as your underlying hardware can go
- ▶▶ Storage media and transport independent architecture
- ▶▶ Full interoperability with all compliant iSCSI initiators





REUVEN M. LERNER

# Google Maps

Google maps out excellent Web services for keen developers.

During the past few months, we have looked at a number of Web services. Web services is a catch-all phrase for the ways in which Internet companies are making their data available to the general public, for use in people's own applications. Thus, Amazon makes its product catalog available for us to create on-line stores and pricing programs, eBay allows us to search through (and bid on) products available for sale, and Google makes its search results available for viewing and manipulation. Each company restricts the ways in which we are allowed to use the provided data, but the trend appears to be toward additional openness and availability.

Sometimes, that openness comes in a package that is slightly different from the standard form of Web services. That is, some companies make their data available using specialized libraries that call the services for you, hiding the specifics of the calls from your application. One of the most famous examples, and the one that we look at this month, is Google Maps. I have found Google Maps to be one of the most compelling and powerful Web applications out there. Not coincidentally, Google Maps was one of the first applications to make use of Ajax, a term that describes how we can use a combination of JavaScript and XML to grab data from remote servers and then use the results to update a Web page dynamically.

This month, I explain how easy it is to create maps using the Google Maps API. We create some basic maps and even put up small markers indicating locations of interest to us. This will serve as a building block to creating our own mashups, the increasingly popular term for the use of Google Maps to display information culled from a separate database.


## Basics

Google Maps, like most Web applications, divides the work between the client (a Web browser) and a server. However, the traditional division of labor has been fairly unequal, putting almost all of the computational onus on the server, giving the browser responsibility for display alone. Ajax changes this, using one or more JavaScript libraries that know how to manipulate the data being displayed in new and interesting ways.

Although Google may someday release an API that will allow us to create our own Ajax applications with its map information, the current release requires that we install and use everything in a single package. That is, Google provides a JavaScript library—or more precisely, a link to a JavaScript library located on Google's servers—that we incorporate into our pages and then use to create a map.

In order to display maps, we need to use that JavaScript library. However, both to keep track of who is using the API and also to ensure that it is being used according to the rules, the library is available only to holders of a key.

Now, we have seen this sort of restriction before, both in Amazon's Web services and also Google's main Web services (that is, for search results). However, the key used in Google



Maps is somewhat different; it is keyed both to a particular person (with a Google account) and to a particular URL. This means a map key that works at <http://www.example.com> will not also work at <http://www.example.net>.

The first step in using the Google Maps API is to decide under which URL you want to put the maps. I decided to create a new Apache virtual host on my system, which I named [maps.lerner.co.il](http://maps.lerner.co.il). I then registered with the Google Maps API page ([www.google.com/apis/maps](http://www.google.com/apis/maps)), indicating that my applications would be under the URL [maps.lerner.co.il](http://maps.lerner.co.il). Several seconds later, I was greeted with a page containing my API key, as well as a simple starter page that can display a map. The key is a very long string of ASCII characters, separated by spaces.

Because we will base our applications on HTML, we should take a close look at it:

```
<!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>
    <script
src="http://maps.google.com/maps?file=api&v=1&key=
ABQIAAAQK9JhAXQ9eq-G55qgu1ExScF-BH9Y-SIKcAjU8YFS_
uTREdFBSs2-11UWY0kXbUv6argoPyrx3YTg"
type="text/javascript"></script>
  </head>
  <body>
    <div id="map" style="width: 500px; height: 400px"></div>
    <script type="text/javascript">


    var map = new GMap(document.getElementById("map"));
    map.addControl(new GSmallMapControl());
    map.centerAndZoom(new GPoint(-122.1419, 37.4419), 4);

    </script>
  </body>
</html>
```

The HTML document begins by declaring its DOCTYPE, which turns out to be strict XHTML. XHTML is a wonderful idea and ensures that HTML is structured according to all of the strict XML rules. That said, many HTML pages do not adhere to this standard and thus are considered either transitional (meaning, XHTML with a liberal eye) or nothing at all. Because Google Maps tries to be compatible with as many browsers as possible, it benefits greatly from strict adherence to XHTML.

In the `<head>` tag, we see that there is a `<script>` tag,

# 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 service 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

which loads JavaScript source from Google's servers at **maps.google.com**. This ensures that the latest version of the JavaScript library is always available to us and our users. Google promises that when it upgrades the map API, it will give a one-month grace period to allow developers to learn about incompatible changes they might have made.

In the body of the HTML document, we then have a `div` tag, whose ID is `map`. This is the node that we will be passing to Google's JavaScript library. The style attribute passed to the `div` tag contains a width and a height; these determine the size of the map. Your site can display any combination of width and height for the map, allowing you to make adjustments for your particular site design.

Inside of the `div`, we finally get to the heart of the matter, with three calls.

First, we create an instance of a `GMap` object. As you might imagine, a `GMap` represents a particular map within the world of Google Maps. We attach the `GMap` object to the node with the ID of `map`. (If the element does not exist, the map will not appear on the screen.) This means, by the way, that you can have more than one map on a particular Web page—simply create multiple `<div>` tags, each with its own unique ID attribute, and attach different instances of `GMap` to each `<div>`.

Once we have created an instance of `GMap`, we can send

around the point described by the `GPoint` and to display the map at level 4. The closest zoom level is 1, and the farthest away is 15, with levels 16–18 showing different types of wrapping. Level 4 allows you to see streets and is a good starting point for people using a map.

It is important to realize that `GPoint` objects are created with longitude and latitude as their parameters, and not the reverse. This is probably because Google engineers thought in terms of `x` and `y` coordinates, which are more natural for math and science people. However, coordinates are often given in latitude, longitude pairs, as opposed to the reverse—so be careful before blindly entering coordinates into a program without checking their order and meaning.

The `GPoint` created in this default document is in Palo Alto, California, presumably pointing to Google's offices. To look at another area on the map, simply substitute a different set of coordinates. For example, to look at Skokie, Illinois (where I'm currently living), I simply substitute a different set of coordinates:

```
map.centerAndZoom(new GPoint(-87.740070, 42.037030), 4);
```

Sure enough, when I reload my page, I'm now looking at a map of Skokie, rather than Palo Alto.

Finally, Google provides us with the ability to switch between three different views, known as `map`, `satellite` and `hybrid`. By default, these controls are shown in the top-right corner and appear thanks to the line:

```
map.addControl(new GMapTypeControl());
```

As you can probably guess, the above line sends an `addControl` message to our map object, handing it a new instance of `GMapTypeControl`.

## Markers

Finally, let's look at how we can create a marker, as it is known, on our map. A marker lets us identify a particular point on a map, showing it with one of the Google Maps icons that users recognize. Moreover, although we see this functionality today, we easily can create JavaScript handlers for our markers—such that clicking on a marker causes a JavaScript function to be executed and presumably change the look of our map somehow.

To create a marker, we create a new instance of `GMarker`, passing it a `GPoint`:

```
var myMarker = new GMarker(new GPoint(-87.740000, 42.030000));
```

Now that we have created our marker, we can display it on the map:

```
map.addOverlay(myMarker);
```

If you add the above two lines within the `<script>` section of the HTML file, you immediately will see a red marker appear on the screen.

Now, here's where some real magic begins. Everything that we have done so far is done in JavaScript and HTML. Both of these are read and handled by the browser, but they are created by the server. This means that if we create our HTML file not as a static file, but rather dynamically (that is, from a server-side program), we can do all sorts of

## Because Google Maps tries to be compatible with as many browsers as possible, it benefits greatly from strict adherence to XHTML.

it messages to control its behavior. For example, we can add a control to it, allowing us to zoom in and out. In this document, for example, we add a small map control by invoking the `addControl()` method, passing it a new instance of `GSmallMapControl`. The `GSmallMapControl` contains `+/-` buttons for zooming, as well as arrow buttons for moving the map without having to drag the mouse.

Google provides two other control types as well, known as `GSmallZoomControl` (which has only the `+/-` zoom buttons) and `GLargeMapControl` (which includes everything that the `GSmallMapControl` does, plus buttons that allow you to jump to a particular zoom level). The controls always appear in the top-left corner of the map, and there is no way to stop you from instantiating more than one of these controls. This means if you aren't careful, you might create more than one control, leading to an ugly map and site.

After creating our map and adding a control to it, we then instruct the map to show us a particular point. Points in a Google map are represented with the `GPoint` data structure, which represents a single point of longitude and latitude. Longitude and latitude can be represented with either degrees or floating-point numbers; for obvious reasons, `GPoints` are constructed using the latter. The example document has the following call:

```
map.centerAndZoom(new GPoint(-122.1419, 37.4419), 4);
```

The above line of JavaScript sends a `centerAndZoom` message to the map object. It instructs the map to center itself

neat things with the JavaScript.

For example, we can create multiple GMarkers, simply by assigning them to different variables and then attaching each of them to the map. If our Google Maps page is being run by PHP, we can write a short PHP program that inserts appropriate JavaScript code into the page. For example:

```
<?php
$a = array(-87.740070, -87.730000);
$count = 0;

foreach ($a as $v) {
    echo "var myMarker$count = new GMarker(new GPoint($v,
    ↪42.037030));\n";
    echo "map.addOverlay(myMarker$count);\n";
    $count++;
}
?>
```

If we put the above inside of the <script> section of our page, and if we then rename the page to index.php (instead of index.html), we quickly will see two markers on the page, with slightly different longitudes and the same latitude.

Notice how the above code uses PHP's echo command to insert text into the HTML document when it executes. Also notice that we need to add semicolons—one to end the line of JavaScript (inside of the quotes) and one to end

the line of PHP (outside of the quotes). These sorts of issues are always a headache when writing a program that writes another program. Finally, notice how we had to create arbitrary new variable names to avoid using the same variable over and over, and thus losing all but one of the markers. The simplest way to do this is with a \$count variable, which then assures that you will have a unique variable name for each marker.

### Conclusion

Google Maps is a wonderful Web application. But for developers, it's also a platform on which we can create all sorts of new applications and services that depend on maps. In particular, by dynamically creating an HTML document from a programming language, we can insert data that JavaScript can then incorporate into a map. Next month, we will see how we can do this, creating our own mashup—grabbing information from one data source and then displaying it on a Google map.■

**Resources for this article:** [www.linuxjournal.com/article/8393](http://www.linuxjournal.com/article/8393).

---

Reuven M. Lerner, a longtime Web/database consultant, is currently a PhD student in Learning Sciences at Northwestern University in Evanston, Illinois. He and his wife recently celebrated the birth of their son Amotz David.

# Hurricane Electric Internet Services...Speed and Reliability That You Can Depend On!

**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



MARCEL GAGNÉ

# If Only You Could Restore Wine

Back up a bit and see how you can keep your files safe.

**François!** What is that truck doing at the back door to the restaurant, and why are they loading up the files from our office? I'm not angry, *mon ami*, and I've already asked them to put everything back. Surely there must be some simple explanation for this. *Quoi?* Ah, I see where the error occurred. François, when I asked you to arrange for an off-site backup of our files, I didn't mean I wanted every piece of paper in the office moved to another location. I was talking about the files on our Linux systems—data, *mon ami*. Yes, I agree, I should have made my request clearer.

What I really wanted was for you to find a simple, easy-to-use backup program that would handle network backups, so that we could store the information from the Linux desktops in this restaurant to some of our off-site servers. Don't worry, François, I've got just the programs to make things very easy for you. In fact, I will be showing them to our guests as soon as they arrive.

But they are already here! Quickly, François, to the wine cellar while I help our guests to their tables. Bring back the 2000 Napa Valley Cabernet Sauvignon we were submitting to, ahem, quality control, earlier today. *Vite!* Please, *mes amis*, sit and make yourselves comfortable.

François and I were discussing backups and backup software. Every Linux system comes with some basic, classic and powerful

backup tools. These include the tar, cpio and rsync commands to name only a few. Many major distributions provide front ends to these commands via their own administrative interfaces. Today, I thought it might be interesting to visit some alternative backup software, all of which make backing up and restoring data easy.

Ah, François, you have returned with the wine. Excellent. Please, pour for our guests.

The first item on our menu is Jean-Remy Falleri's Keep, a simple, easy-to-use backup utility for KDE. What makes it attractive for a desktop user is that Keep can sit quietly in the background and regularly back up your directories. It does this by using its own KDE service daemon. Backups also can be run at any time with the click of a button, and restores are a piece of cake. Because the package uses rdiff-backup to do its work, you need to install this as well. Source is available from the Keep Web site (see the on-line Resources), but I found binaries easy to get from various contrib sites as well.

Once Keep and its rdiff-backup prerequisite are installed, start Keep by running the command, `keep`. A simple window with five icons appears (Figure 1). The top three options are the ones you will use most often: adding a directory to the backup list, restoring from a backup or running a backup right now. Existing backup lists can be edited, and Keep also provides a simple backup log.

When Keep runs for the first time, you'll see a message at the bottom of the main window indicating that the Keep daemon is not running. That's because the daemon starts when you start KDE. You can click the Load button to start it, but if you run into problems, have no fear. You also can fire up the daemon via the KDE Control Center (command name, `kcontrol`). Look under KDE Components, and start the daemon from the Services Manager.

To create a backup, click the Add directory to backup button. This starts a wizard-like dialog that begins by asking you to select a directory for backup (for example, your home directory). Click Next, and a KDE file selector appears from which you can identify the location of your backups. Click Next again, and you are almost done. The final screen is where you select the backup interval (Figure 2). The default is to run automatically every three days and to delete archives after 60 days, but you can change this to whatever you like. Extra options let you select compression and whether or not so-called special files are excluded. There's also a check box to fine-tune what exactly gets backed up using the Advanced Configuration dialog.

Click Finish, and your backup definition is created and scheduled to run at a later time. When the Add a backup window closes, you'll be back at the main Keep window. Feel free to create more than one backup definition with different backup intervals. At the bottom of the main window is a View backup log button where you can check on the status of your current backup. If you don't want to wait for your scheduled backup to run, click the Backup now button. A small window appears




Figure 1. Keep's interface is compact and easy to use.

# 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 ClusterWare™ Enterprise *Web Suite*. Its centrally-managed, highly available architecture makes large pools of Linux servers act like a single, consistent 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



## REALLIFELINUX

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

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



Penguin Computing and the Penguin Computing logo are registered trademarks of Penguin Computing Inc. Scyld ClusterWare 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.



Figure 2. Keep runs automatically scheduled backups.

with the various backup jobs you have defined. Click on the one you like, and then click OK to launch the backup.

As the days go on, Keep starts to build a list of incremental backups based on the original full backup. Each of these snapshots are dated in preparation for the inevitable. After all, the whole point of a backup is to prepare for the day when you need to restore something that has gone missing through some disaster, *non*? To restore a file or folder, click the Restore a backup button. Choose a backup directory from the list presented, click Next, and then choose where you want your directories restored—either in the original directory or in an alternate location of your choosing. If you want to restore an individual file, you’ll want to restore to an alternate directory rather than overwriting your own.

Click Next, and a list of dates appears. This is how you decide to which point in time you want to return. Click on a date, then click Finish, and the restore process begins. Keep is decidedly simple, but if your backup needs are equally simple, this is a program worth looking into.

For the GNOME users out there, we have Aigars Mahinovs’ appropriately named Simple Backup. This program was created as part of Google’s “Summer of Code” and was envisioned as an Ubuntu application. Unlike Keep, Simple Backup is a two-part application, with the backup configuration as one application and the restore as the other. Like Keep, Simple Backup runs predefined backups in the background according to whatever schedule you assign. In many ways, however, Simple Backup is much more flexible and powerful. You can get Simple Backup from SourceForge (see Resources).

The first step in using Simple Backup is to start the configuration program. This is done by selecting Simple Backup Config from the GNOME System menu in the top panel. Because this qualifies as an administrative task, you’ll need to enter the root password to proceed (or your password if you are running Ubuntu). This brings up the Backup Properties dialog (Figure 3).

Three radio buttons allow you to select your backup settings. By default, Simple Backup does standard daily and incremental backups of user data to the `/var/backup` directory. Large data




Figure 3. The first step in creating a Simple Backup is to define a backup configuration.



Figure 4. To restore from a Simple Backup, select the date of the backup you want and choose the directories you want restored.

files, greater than 100MB, are excluded as are multimedia and temporary files.

The Include and Exclude tabs allow you to specify which directories or files you want to have backed up. On desktop systems, I tend to back up my data, the system configuration files and nothing else. With servers, I back up everything. Your choice may likely be somewhere in between. The most interesting option here is the Destination tab, and the reason I suggested that you choose a custom backup configuration on the General tab. You still have the option of choosing the default backup directory of `/var/backup` as well as an alternate directory.

To recover a directory using Simple Backup, click System on the GNOME top panel, and select Simple Backup Restore from the Administration submenu. Once again, you’ll be asked for a confirmation password, after which the Restore files/directories dialog appears (Figure 4).

Your default backup location (or restore source) is indicated at the top of the window; however, if you have backups in a different location, click the Custom restore source check box, then enter the pathname in the location field and click Apply. A list of available backups appears in the drop-down window below. Click on one, and the folders from which you can restore will show up in the main central area. Navigate to the folder you want (by clicking the arrows to expand subdirectories), and make your selection.


You now have two choices. The first is to restore the folder as it was, in its original location. In some cases, the right choice will be Restore As, which lets you select an alternate

**What makes it attractive for a desktop user is that Keep can sit quietly in the background and regularly back up your directories.**

location or name for the directory you are restoring. When you have made your choice, a confirmation box asks you whether you really, really want to restore the folder to the location specified. Assuming the answer is yes, click Yes. The whole process from backup to restore is very simple. If I could make one recommendation, however, it would be to provide a log progress window and an easily accessible log. Otherwise, Simple Backup is very much as the name would indicate.

Finally, I would like to show you Johnathan K. Burchill's KDar, or KDE Disk Archiver, a friendly, graphical interface to Denis Corbin's powerful command-line dar utility. Of the programs on today's menu, this is by far the most flexible, for reasons I'll explain shortly. Where the other programs work at the directory level, KDar can restore individual files as well. It can do full and incremental backups, and it can break up the archives into slices to fit on the storage media with which you choose to work. This media can be a CD-ROM, DVD and so on. You can get KDar from SourceForge (see Resources) where source bundles are available. Should you prefer binary packages, KDar is easily found on a number of contrib sites.

Once the package is installed, fire up KDar by running the kdar command (use the Alt-F2 quick launch if you prefer). A splash screen flashes a moment before the actual interface starts (you can turn off the splash screen in the configuration dialog under Settings). When the program starts, the main window looks fairly plain (Figure 5). Along the top is a pretty standard menu bar with some quick access icons directly below.



Pause your mouse cursor over the icons, and tooltips identify them for you. Below that are two large, empty panes. The top pane lists archives and files, and the bottom is a log window.

Figure 5. KDar's Interface at Start Time

The easiest way to create a backup is to click on the Create icon or select Create from the Archive menu off the menu bar.



## 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™






Figure 6. KDar's backups provide a graphical progress report with estimated time to completion.

You'll be asked for the top-level folder you want to back up, where you want the backup stored and so on. You'll also be asked to configure creation options, and these are substantial, so take some time to look them over. These options include compression, cryptography, files and folders to include or exclude, and file types to exclude (such as MP3, AVI and so forth) and much more. Once you have finished with the creation of your backup, you can save the profile, export the dar command (more on this in a moment), do a dry run of the backup (without actually writing) or simply start the backup. A progress window lets you know how things are progressing (Figure 6).

When it comes time to restoring a file or folder, your first step is to pull up an existing archive. Click File on the menu bar and select using the Open menu (or go for Open Recent). The top pane of the KDar main window displays the archive with a small arrow beside it. Click the arrow to expand the folder list, and each subdirectory also opens with an arrow beside it until you get to the file level (Figure 7).

Select the file, directory or combination of both that you want restored, then right-click on your selection. This brings up a small menu from which you can choose to restore, do a diff comparison of your backup against the current files or do a test restore (Figure 8). All of these choices also are available by clicking Archive on the menu bar.

KDar provides a log of the restore process, including the number of files restored, the time taken and any errors that were encountered.

Whenever you run a backup or a restore using KDar, there is a button on the final screen that lets you export the equivalent dar command to a bash shell script. This is important, because it makes it easy to create cron jobs for your backups. Although KDar is much more flexible than any of the other candidates I've covered, it lacks a daemon that backs up in the background. Nevertheless, the shell scripts it generates makes KDar (and its dar counterpart) suited to more complex environments.

Ah, *mes amis*, if only there was some way to restore the clock to an earlier time. It seems that this is still somewhat beyond the talents of even the most skilled programmer. Until such a time as this wondrous package becomes available, I'm sure that François will not mind if we keep the restaurant open just a little longer so he can refill your glasses once more before the final "Au revoir". It's also too




Figure 7. KDar can restore individual files as well as directories.





Figure 8. The final step in a KDar restore allows you to store or load profiles. You also can export the commands to a shell script.

bad that we can't restore all this wine. Can you imagine it, *mes amis*? And endless wine cellar. That dream is very much alive. On that note, please raise your glasses, *mes amis*, and let us all drink to one another's health. *A votre santé! Bon appétit!* ■

**Resources for this article:** [www.linuxjournal.com/article/8940](http://www.linuxjournal.com/article/8940).

Marcel Gagné is an award-winning writer living in Mississauga, Ontario. He is the author of the all new *Moving to Ubuntu Linux*, his fifth book from Addison-Wesley. He also makes regular television appearances as Call for Help's Linux guy. Marcel is also a pilot, a past Top-40 disc jockey, writes science fiction and fantasy, and folds a mean Origami T-Rex. He can be reached via e-mail at [maggagne@salmar.com](mailto:maggagne@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).

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.



**ZT Opteron Server A8125**

**2 x Dual-Core AMD Opteron™ Processors 265**  
(2 x 1MB L2 Cache, 64/32 Bit)

- AMD® 8131 + 8111 Chipset Server Board
- 1GB DDR 400MHz ECC/ Reg. SDRAM (Up to 16GB)
- 8 x 250GB 7,200rpm SATAII Hard Drive (SAS / SATA Convertible)
- 8 x 1" SAS / SATA Hot-swappable Drive Bays
- Adaptec® 8 Channel SATAII 64bit 133MHz RAID Controller (RAID 6 Ready)
- DVD-RW and Floppy Drive
- 2 x LAN with Broadcom 5704 Gigabit Ethernet
- 4U Rackmount Chassis w/ 650W Redundant-Cooling Power Supply
- SuperDoctor III Server Management Software
- 3-Year Limited Warranty + 1-Year 4 Hour Onsite Warranty

**\$3,499**

**ZT Super Power Opteron Server A8126**

**2 x Dual-Core AMD Opteron™ Processors 280**  
(2 x 1MB L2 Cache, 64/32 Bit)

- AMD® 8131 + 8111 Chipset Server Board
- 2GB DDR 400MHz ECC/ Reg. SDRAM (Up to 16GB)
- 6 x 500GB 7,200rpm SATAII Enterprise Hard Drive (SAS / SATA Convertible)
- 6 x 1" SAS / SATA Hot-swappable Drive Bays
- Adaptec® 8 Channel SATAII 64bit 133MHz RAID Controller (RAID 6 Ready)
- Slim CD-ROM and Floppy Drive
- 2 x LAN with Broadcom 5704 Gigabit Ethernet
- 2U Rackmount Chassis w/ 500W Redundant Power Supply
- SuperDoctor III Server Management Software
- 3-Year Limited Warranty + 1-Year 4 Hour Onsite Warranty

**\$6,999**

Find out how ZT can help maximize your business solutions

Go to  
Call

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

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

Promote Code # LJ0606



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



DAVE TAYLOR

# Coping with Aces

An Ace in the hole helps this longest *Blackjack* exercise ever.

Somehow, writing this *Blackjack* game is starting to feel like the programmatic equivalent of that *Three Stooges* skit where “slowly he turned, step by step...”, but we’re still going to have to work on the core logic of the game before we’re ready to write the fun interface elements.

This month, in fact, we might well find that we have to tear some of the earlier script apart and rebuild it to compensate for a troubling aspect of the game of *Blackjack*: an Ace can be either high or low, which is to say that it can be worth one or 11 points. Dealt two aces, you then have a number of different possible values, and that’s a problem.

## The First Ace Is Always Worth Eleven

It turns out that there’s a sneaky way you can solve this problem simply by maximizing the value of the first Ace encountered, as long as the overall value of the hand doesn’t exceed our cap of 21 points. So, two Aces would be worth 11 + 1 automatically (the first is maximized, but the second is not because it would push us over 21 points).

The portion of the code that must be rewritten to compensate for this Ace valuation strategy is the `handValue` function:

```
function handValue
{
    # feed this as many cards as are in the hand
    handvalue=0 # initialize
    for cardvalue
    do
        rankvalue=$(( $cardvalue % 13 ))
        case $rankvalue in
            0|11|12 ) rankvalue=10 ;;
            1 ) rankvalue=11 ;;
        esac
    done
}
```

A lot of programmers talk about highly efficient code as being “elegant”, but in my experience, most elegant code is also lazy code.



```
    handvalue=$(( $handvalue + $rankvalue ))
done
}
```

This is the “before” picture from last month. Notice that the second line in the case statement currently assigns a rank value of 11 to every Ace encountered. Clearly that’s a bug!

To change it, however, I need to add a new variable that keeps track of whether I’ve already seen a previous Ace in the hand. I ingeniously call that `seenAce`:

```
function handValue
{
    # feed this as many cards as are in the hand
    handvalue=0 # initialize
    seenAce=0
    for cardvalue
    do
        rankvalue=$(( $cardvalue % 13 ))
        case $rankvalue in
            0|11|12 ) rankvalue=10 ;;
            1 ) if [ $seenAce -eq 1 ] ; then
                rankvalue=1
            else
                rankvalue=11 ; seenAce=1
            fi ;;
        esac

        handvalue=$(( $handvalue + $rankvalue ))
    done
}
```

Looks like it’ll do the job—or will it?

The problem here is best illustrated with a hand like 9 + 10 + A. That’s a valid *Blackjack* hand and should be worth 20 points. But `handValue` will score it as 30 points, and the program will incorrectly classify that hand as a bust.

Solving this isn’t too hard once the problem is recognized, but that’s the great challenge of writing any code, isn’t it? To anticipate and characterize bugs and glitches properly. The solution is often quite simple, but knowing there’s a bug in the first place, ah, that’s where the great programmers find their calling!

The solution in this situation is that we need to deduct ten points from the hand score if it’s more than 21 points and there’s an Ace—a condition that turns out to be added easily to the tail end of the function:

```
    handvalue=$(( $handvalue + $rankvalue ))
done

if [ $handvalue -gt 21 -a $seenAce -eq 1 ] ; then
    handvalue=$(( $handvalue - 10 ))
fi
}
```

# PGI<sup>®</sup> Compilers are building the 64-bit applications infrastructure.

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

Visit [www.pgroup.com](http://www.pgroup.com) to learn what PGI Compilers and Tools can do for you.



**The Portland Group**<sup>™</sup>  
[www.pgroup.com](http://www.pgroup.com) ++ 01 (503) 682-2806

This is the first time I've used a complex conditional statement in our script, but you're already familiar with this type of multi-expression conditional. If we were using a C-like language, the conditional might look like:

```
if ( ( handvalue > 21 ) && (seenAce == 1))
```

The snippet in the shell script shown above is the equivalent conditional, with the `-a` serving as the logical AND statement. It wouldn't work in this context, but `-o` is the logical OR statement in a shell test conditional too, and if you need to, you can use parentheses for grouping.

To test our new code, I'm going to replace the main body of the program temporarily with a few preloaded test hands and see what kind of hand values are returned:

```
echo "Starting out with two aces..."
handValue 1 14
echo "handvalue = $handvalue"

echo "now testing 9 + 10 + A"
handValue 9 10 1
echo "handvalue = $handvalue"

echo "and, for good luck, testing K + A"
handValue 12 1
echo "handvalue = $handvalue"
```

First, I'll run this with the original `handValue` function, anticipating mistakes:

```
Starting out with two aces...
```

```
handvalue = 22
now testing 9 + 10 + A
handvalue = 30
and, for good luck, testing K + A
handvalue = 21
```

Yup. That's not good. We'd be quickly run out of Vegas for that sort of counting.

Now, I'll slip in the new `seenAce` code segments explained earlier and try this same set of test hands:

```
Starting out with two aces...
```

```
handvalue = 12
now testing 9 + 10 + A
handvalue = 20
and, for good luck, testing K + A
handvalue = 21
```

What do you know, it looks like we've come up with a savvy way to allow the Ace to have two possible values without a major rewrite of the code.

### The Virtue of Lazy Coders

Good. Indeed, it's my belief that the best programmers are actually lazy and want to solve problems in the easiest and most efficient way possible. Laziness breeds ingenuity, remember, so although I could have rewritten the blackjack script to use an array of possible hand values to model the multivalued hand, why bother? The fact that a given hand has more than one value isn't really important as long as we can compensate for that fact correctly in the code.

A lot of programmers talk about highly efficient code as being "elegant", but in my experience, most elegant code is also lazy code. I know that I'm constantly looking for those smart shortcuts, those insights that let me create something that might be less efficient in its performance, but far easier to code, far faster to debug and far speedier to deploy in the field.

One great skill that programmers can nurture is being able to recognize quickly the good-enough solution too. Highly analytic by nature, we code geeks suffer from a little bit of perfectionism, and writing the perfect routine at the cost of additional days or weeks of development easily can end up being less utilitarian and less useful than having a pretty decent routine that does the job and can be improved later, in the next release, a maintenance patch or whatever.

Is this laziness what causes us to have software with so darn many bugs though? I don't think so. I think bugs in products are due to the ever-increasing level of complexity of software, be it an administrative tool for a Linux box, an Apache module or an Ajax-y Web-based utility. And software like an operating system or kernel? Of course it's going to have bugs. It's far too complex ever to test for all possible conditions, cases and situations. In fact, seeking efficient solutions that can be pushed out into the field can help reduce bugs. It's not testing software that finds the most egregious problems, but customers putting software through real-world tasks.

I'm not advocating that we should ship sloppy code, however. Simply that in the classic model of alpha and beta releases, getting code into the field ultimately can produce far more robust applications than having it stay in development forever as more and more complex test cases and usage scenarios are pushed through simulators.

But, ahem, I digress!

For now, we've come up with a nice, simple solution to the dual-value problem with Aces, and let's leave our script here for this month. Next month, we'll reintegrate the new code into the main game and add some additional code to detect when either the player or dealer has a blackjack (a two-card 21).■

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).

## We Know Linux!

### Your Embedded Linux Partner

*Supporting Embedded Linux For Over 10 Years*

Single Board Computers	2.6 Kernel	Real-Time
Custom Drivers		Starter Kits
Flash Disk		Enclosures
USB Support		Panel PCs
Fanless CPUs		X-Windows
PCMCIA		CAN
PC/104		HTTP Server
IP Router/Firewall		Data Acquisition

**EMAC can fulfill your Embedded Linux needs, from credit card size systems to rack mounts.**

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)

# Polywell PC Appliance Solutions

## OEM Service for Small to Medium Size Projects

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

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

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

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



1114AIS-2050M, 285+, 32GB, 2TB **\$9,999**  
 1114AIS-2050M, 260HE, 16GB, 1TB **\$4,499**  
 1112ES-2200A, 265+, 2GB, 500GB **\$1,999**



4024AIS 4U 12TB SAS/SATA **\$12,500**  
 4024AIS 4U 6TB StorageServer **\$7,950**

### Linux Appliance PCs

- Custom Made Odd Size Chassis
- AC or DC Power Supply
- Low Power Voltage AMD Sempron™ Processor
- or High Performance AMD Athlon™ 64 Processor
- Diskless or Flash OS Boot Drive
- Swappable Hard Drive, CD-ROM, FDD
- Integrated Graphics, Ethernet, USB
- Optional LCD LED Control Module
- IS2 Audio, MPEG2/4 Hardware Video
- Up to 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,



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

### 64GB DDR 4-way Workstation

- 2 x AMD® Opteron™ Dual-Core Processors 285+
- Up to 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 12-Bay Tower + 600W 80Plus P/S
- 2x250G 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



64G RAM, 2x285+, QuadroFX4500 **\$22,999**  
 32G RAM, 2x265+, QuadroFX3450 **\$7,599**  
 16G RAM, 2x246+, QuadroFX1300 **\$3,750**

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

### 1U Power Saving ISP Server

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



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

Pick the chassis to fit your application

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

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



5U 16-way 865, 32GB, 2TB, 8801T5U **\$19,999**  
 2U 8-way 865, 8GB, 900GB, 8422C **\$6,999**

### Mainstream Linux Systems

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



939NV3000TD12LX starts at **\$399**

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



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

**888.765.9686**

[www.Polywell.com/us/LJ](http://www.Polywell.com/us/LJ)



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

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



MICK BAUER

# Security Features in Red Hat Enterprise 4

Red Hat Enterprise Server proves that less can be more, especially with the help of well-implemented SELinux.

**This month,** I conclude a three-part series on distribution-specific security features. I began with SUSE Linux 10.0, continued on to Debian GNU/Linux 3.1 and this month I discuss Red Hat Enterprise Linux (RHEL) ES 4.

Red Hat Enterprise Linux is a general-purpose Linux distribution targeted to both desktop and server markets. As the name implies, RHEL is intended to be highly robust, stable and scalable; in other words, suitable for production use across large enterprises. And, sure enough, RHEL enjoys the reputation of delivering on all counts. Like SUSE, RHEL even runs on IBM eServer z-Series mainframes.

To a much greater degree than Debian, however, and to a significantly greater degree than SUSE, Red Hat adheres to a strict philosophy of less is more where bundled software packages are concerned. Whereas Debian is composed of more than 15,000 packages and SUSE of more than 4,000, RHEL ES 4 weighs in at a mere 1,730 (if you include RHEL Application Server and Extras packages, which aren't part of the base OS, strictly speaking).

I don't think it's at all euphemistic to say that this is an easily defended design choice. By limiting the number of packages it officially supports, Red Hat has a much smaller attack surface (not to mention help-desk surface). Fewer packages mean less complexity; less complexity means better stability and security (at least in theory).

The downside of this design philosophy is obvious. It means fewer choices in any given tool space (network servers/dæmons, encryption tools and so on), less flexibility and greater likelihood that you'll end up installing third-party packages or even compiling them yourself from source.

As I've said many times in this column, there's no harm in rolling your own, especially when that means you're compiling out (excluding) unnecessary or potentially insecure features. But, nothing beats distribution-supported binary packages when it comes to automated security updates. And, none of the major Linux distributions besides Gentoo has any automated means of applying security patches directly from source code to locally compiled software.

Furthermore, as I'm about to show, RHEL ES 4 is particularly thin in the specific realms of security-enhancing software (with the sole exception of SELinux) and security-scanning software. This doesn't mean that I think RHEL is insecure; its smaller attack surface and its excellent SELinux support are both highly significant. It does mean, however, that you've got fewer choices in *how* you secure your RHEL-based server or desktop system, and even fewer choices in how you use it in security applications, than is the case with other major distributions.

## Installing RHEL ES 4

Red Hat Enterprise Linux ES 4 has a very easy-to-use installation

GUI that, besides installing the base operating system, allows you to select additional software packages, set the root password, set up networking, enable a simple local firewall policy and enable SELinux. After the first reboot, this installer runs additional modules for setting up a Red Hat Network subscription, creating the first nonroot user account and configuring the X Window System.

Personally, I don't care for the Red Hat installer's software package selection module at all. First, it allows you to select only from a subset of the packages available on the installation medium. (That is, as far as I could tell—it could simply be that a few packages I knew were available but couldn't find, such as `gnupg`, were simply buried within categories in which I didn't think to look.) For the packages it does display, the installer shows neither detailed descriptions nor even approximate disk space required.

In addition, its dependency-checking functionality is decidedly primitive. If the software installer can't find something it needs, it returns an error but doesn't give you any means of solving the problem (locating the missing package, deselecting or uninstalling the package with the unmet dependency and so on). Although simplicity may be a virtue, this limited functionality doesn't compare very favorably at all with Debian's `aptitude` package management tool or SUSE's `YaST`. If you want to run this installer module again after installation is complete, it's located in GNOME's Application menu under System Settings under Add/Remove Applications, though I think you might be much happier performing additional software installations using `up2date` or even good old RPM.

So, what security-related packages *are* available in RHEL ES 4? Table 1 lists most of them. In a nutshell, if you want to secure the local system, SELinux and your local firewall policy are very nearly the only tools available to you. If you want to audit and analyze the security of other systems, RHEL ES 4 has very little to offer besides `Nmap`.

On the face of it, this is a decent list of applications; these are all important security-enhancing tools. Notably absent, however, are:

- Any sort of file-integrity checker, such as `Tripwire` or `AIDE`.
- `Syslog-NG`, a much more powerful system logger than the archaic `syslogd` on which RHEL still relies.
- Any sort of virtualization environment (user-mode Linux, `Bochs`, `Xen` and so on).
- The ubiquitous intrusion detection system and packet-logger `Snort`.
- Web security tools such as `squidguard`, `mod_security` and so on.



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)

Table 1. Some Security-Enhancing Packages in RHEL ES 4

Package Name	Description
bind-chroot	Configures your BIND-based DNS server to run in a chroot jail.
dovecot	IMAP server (mail delivery agent) designed for security.
freeradius	RADIUS authentication service for network devices.
krb5-server	Kerberos authentication/encryption server.
splint	Tool for auditing C code for bugs, including security vulnerabilities.
vsftpd	Very Secure FTP Daemon: RHEL's only FTP server, but an excellent choice.
cryptsetup	Tool for creating encrypted filesystems (as virtual block devices).
ethereal, tcpdump	Classic protocol analyzers (aka packet sniffers).
gnupg	GnuPG e-mail/general-purpose encryption tool.
ipsec-tools	Utilities for building IPsec VPN tunnels.
nc	Netcat, a versatile IP packet redirector.
nmap, nmap-front end.	The Nmap port scanner and its GUI front end.
openldap-clients, openldap-servers	OpenLDAP directory and authentication system.
openssh	The most popular free Secure Shell daemon and client.
openssl	General-purpose SSL/TLS cryptographic library and tools.
policycoreutils, setools, setools-gui	Tools for configuring and managing SELinux policies.
selinux-doc	Not installed by default, but you'll want this collection of SELinux documents.
postfix-pflogsumm	Log summarizer for the Postfix mail transfer agent.
spamassassin	Popular SPAM/UCE filter.
stunnel	General-purpose SSL/TLS wrapper for TCP applications.
sudo, usermode	Tools for allowing nonprivileged users to run processes as root.
tcp_wrappers	Provides simple IP-based access controls to TCP applications.
up2date, up2date-gnome	Red Hat's automated network-based software update tool.

You're perfectly free, of course, to download and compile the source code of any of these tools manually. But, you won't be able to leverage up2date's automatic update features on such packages.

So, both in terms of available security packages and the software installer itself, RHEL is a bit underwhelming. On the plus side, I do like the Red Hat Enterprise Linux installer's fire-




Figure 1. Firewall and SELinux Configuration in RHEL ES 4

wall/SELinux module (Figure 1). Both the firewall and SELinux functionality are enabled by default, and the help window on the left-hand portion of the screen explains both settings in plain language.

If you're completely new to SELinux, you can select a warn setting that causes the kernel to log events that violate the local SELinux policy without actually blocking those events. By default, however, SELinux is set to active, using a default policy that restricts the behavior of Apache (httpd), bind, NIS (ypbind), dhcpd, mysqld, ntpd, portmap, postgresql, snmpd, squid and syslogd.

The last thing worth noting about the Red Hat Enterprise Linux ES 4 installer is that both during initial setup, when you enter the root password, and after the first reboot, when you create the first nonroot user account, the installer performs no password complexity checks of any kind (of the sort SUSE's installer performs). It doesn't even warn against choosing an overly simple password via a simple text box like Debian does.

This is unfortunate. Password guessing and brute-force attacks are still very much with us. I was pleased to see, however, that by default, the XScreenSaver utility is configured to lock X sessions by password automatically any time the screensaver kicks in. (If only those passwords that protect XScreenSaver were required to include some combination of mixed upper-/lowercase, punctuation and numerals, I'd be happier still!)

### Automated Updates with up2date

Keeping your system up to date with the latest security patches is absolutely essential on any Linux system. Red Hat was a pioneer in offering automatic updates when it introduced the combination of the up2date utility and the Red Hat Network service offering several years ago, and this system is even more mature now.

The way this works is that when you set up your Red Hat system (any current version), after the first reboot you're prompted to configure a Red Hat Network subscription. A subscription with an RHN Update entitlement is included with every Red Hat product. When prompted, you simply enter the user name and password you'd like to use (one account can be used to manage multiple systems under the same subscription), and then the subscription number printed on the

Activate Your Subscription card that came with your Red Hat installation media.

The net effect of all this (no pun intended) is that you now will have an active subscription to the Red Hat Network service, with a system profile corresponding to your new Red Hat system, which in turn is associated with an RHN Update entitlement that allows your system to check for and download the latest versions of all software packages that are part of the version of RHEL you purchased.

The simplest way to check for and apply security updates is to right-click the icon for the Red Hat Network Alert Notification Tool on your GNOME desktop (it's a glowing red exclamation point if your system isn't up to date, or a blue check mark if it is), and select Check for updates, run up2date and so on, as needed.

You can set up automatic updates by logging on to the Red Hat Network Web site ([www.redhat.com/en\\_us/USA/rhn](http://www.redhat.com/en_us/USA/rhn) for US users) with your RHN credentials, clicking on the Systems tab, clicking on your system's profile, clicking Properties and checking the box next to Automatic application of relevant errata (Figure 2). Obviously, you shouldn't enable this feature on high-availability or change-controlled systems, because software patches always have the potential to introduce other bugs or conflicts.




Figure 2. Configuring Auto-Updates Via the RHN Web Site

Although the up2date/RHN system is mature and feature-rich (especially for large organizations with the need and ability to pay for network management and provisioning entitlements), as a Linux desktop user, I find it more difficult to use than Debian's apt system (which is more primitive in some ways, but easier to script) or SUSE's YaST Online Update system (which is much easier to configure).

Oddly, as with many other aspects of RHEL, up2date configuration options appear to be spread across multiple GUIs, including the Red Hat Network Web site, unless of course you configure things from a shell (in which case everything you need is in `/etc/sysconfig`). If you administer Red Hat on servers (that may not even have the X Window System installed, which is always a good policy on hardened systems) or are otherwise command-line-centric, I'm sure up2date and other Red Hat functions are easy to learn. Ironically, I find many of RHEL's GUIs, which are, of course, supposed to simplify things, confusing. (But maybe it's just me!)

## SELinux on Red Hat Enterprise Linux

As we've seen, RHEL seems to rely very heavily on SELinux for system security. This is hardly a sloppy or mentally lazy design choice; SELinux provides a comprehensive and granular array of mandatory access controls against system users, applications, processes and files. As described in the previous section, the included targeted SELinux policy provides default controls on some of the most commonly used applications.

This default policy's behavior can be tweaked easily using the Security Level applet accessible via GNOME's Application→System Settings menu (Figure 3). The same applet can be used to configure a simple local firewall policy.

The implementation of SELinux in RHEL ES 4 is truly commendable for its simplicity, not to mention the very fact that it's enabled by default. That's the good news; the less-good news is that to create a custom SELinux policy, that is, one that uses tighter or looser controls than the included policy or one that addresses other applications, you're going to have to do some reading. The best place to start is the Red Hat Enterprise Linux 4 Red Hat SELinux Guide, available at [www.redhat.com/docs/manuals/enterprise/RHEL-4-Manual/selinux-guide](http://www.redhat.com/docs/manuals/enterprise/RHEL-4-Manual/selinux-guide).

You'll also probably want to install some extra GUI tools for this purpose too, namely the setools and setools-gui packages. These packages provide, among other things, sepcut, apol, seaudit and seuserx. For more information on what these

## 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**




Figure 3. Security Level Configuration Applet

tools do and how to use them, see the documents in `/usr/share/doc/setools-1.5.1` (the directory name on your system may reflect a different version number).

### Firewall Functionality in RHEL

I've already mentioned the Security Level applet in RHEL ES 4's GNOME desktop. Unlike with SELinux, this applet doesn't give you much more in the way of configuration options for the local firewall policy than you get at installation time. This policy allows all outbound network transactions (originating from the local system), and blocks all inbound network transactions (destined for the local system) except the services you select here. Those services are, as in the Red Hat installer, HTTP, FTP, Telnet, mail (SMTP) and SSH.

In the Security Level applet, you also can specify a list of other ports in the form `[port #]:[protocol]`, for example `689:tcp, 53:udp, 53:tcp`. If you need anything fancier than that, you have to write your own iptables statements

from scratch. Happily, you can do so simply by adding or editing lines in the file `/etc/sysconfig/iptables`. See the iptables man page and the Red Hat Enterprise Linux 4 Security Guide for more information.

### Directory Services and PKI

It's worth mentioning that Red Hat recently acquired Netscape Directory Server, and has updated it and rebranded it as Red Hat Directory Server. This is being positioned as a commercially supported alternative to OpenLDAP or Sun Java System Directory Server. Although not included with RHEL (it's an add-on product that costs extra), it's worth mentioning as a key component of Red Hat's security vision. RHEL does include fully supported OpenLDAP packages, however.

In the same vein, Red Hat Certificate System provides a commercially supported PKI solution. It too is an add-on product not included with RHEL. OpenSSL is included with RHEL, of course, but without any additional setup tools such as TinyCA.

### Conclusion

I have mixed feelings about Red Hat Enterprise ES 4's security features. On the one hand, RHEL doesn't offer anywhere near as many different security-enhancing software tools as Debian GNU/Linux or SUSE Linux. Entire categories of security tools that are well represented in other major Linux distributions (integrity checkers, intrusion detection systems, virtualization environments and so on) are absent.

On the other hand, Red Hat has clearly maintained an unparalleled level of control over the size and scope of its distribution. It has made hard choices about what it will support and maintain, and what it will not, which surely reduces the attack surface of Red Hat systems. I have no doubt that Red Hat's security team has an easier time responding to vulnerabilities in RHEL's 1,730 packages than the Debian security team does with that distribution's 15,000-plus packages.

Furthermore, by not only including SELinux in RHEL 4 but also enabling it by default, Red Hat has taken a very bold step. The kernel-level mandatory access controls provided by SELinux provide the potential to mitigate many of the risks one might otherwise use add-on utilities to address. Furthermore, because this sort of technology is proactive, designed to prevent bad behavior, it's inherently stronger than intrusion detection, integrity checking and other reactive technologies (though it would be better if RHEL had *both* proactive and reactive measures—even with SELinux, bad things can happen).

Whether you find RHEL to be lean and mean or limited and clunky will depend on your particular Linux needs. I'll allow that some of the reasons I'm not as keen on RHEL as I am on Debian and SUSE are specific to my job as a security architect and consultant. I rely on a specialized set of tools, most of which RHEL has judged to be unnecessary for its target market—presumably IT professionals in corporate settings. Still, it seems to me that if I needed to secure a corporate Web server running RHEL, with or without SELinux, I'd still want to install `mod_security`, `Squidguard`, `Syslog-NG` and other tools manually that RHEL presently lacks. ■

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".

## Cost-efficient performance for high performance computing

Intel® Xeon™ Processor powering ServersDirect Systems provides outstanding price/performance solutions for your HPS needs



**\$999**

### SDR-1108T 1U Xeon™ Simple and Affordable Server

Entry-level server solutions for networking, shared Internet access, file and print sharing

- 1U rack designed with 2 x 1" Hot-swap SATA Drive Bays
- Intel® Xeon™ 2.8 Ghz with 800MHz 2MB Cache (Dual Processor Option)
- Intel® E7320 (Lindenhurst VS) Chipset
- Kingston 1024MB DDRII 400 ECC Reg.memory
- 1pcx120GB SATA 7200RPM hard drive
- ATI Rage XL SVGA PCI video controller with 8 MB of video memory
- 6300ESB (Hance Rapids) SATA Controller (2x Drive support)
- RAID 0, and 1 support
- 2x Intel® 82541GI Gigabit Ethernet Controllers

### 2U Xeon™ Applications Server SDR-2116T

Optimal solution for mail and messaging, databases, and applications



- Intel Xeon™ 2.8 Ghz with 800MHz 2MB Cache (Dual Processor Option)
- 2U Chassis with 460W Power Supply
- Kingston 1024MB DDRII 400 ECC Reg.memory
- Western Digital 250GB SATA RAID Drive
- 6 x 1" Hot-swap SATA drive bays
- Integrated ATI Rage XL SVGA PCI video controller
- One Intel® 82541PI Gigabit Ethernet controller
- Dual-channel SATA controller supporting RAID 0 and 1

**\$1199**

### 3U Xeon™ Storage Server SDR-3309T

Provides optimal solution for Databases and applications



- Intel® Xeon™ 2.8Ghz with 800FSB 2MB Cache (Dual Processor Option)
- 3U Chassis with 650W Redundant Power Supply
- Intel® server board with Intel® E7520 (Lindenhurst) Chipset
- Kingston 1024MB DDR333 ECC Reg. (2x512MB)
- 1pc x Western Digital 250GB SATA RAID Drive
- 16 x 1" Hot-swap SATA drive bays
- Integrated ATI Rage XL SVGA PCI video controller
- Dual Intel® Gigabit Ethernet Controllers
- 2 SATA Channels with RAID 0 and 1

**\$2,699**

### 4U Enterprise Server SDR-4305T

Easily scalable storage solution with Hot-swap functionality for growing business



- Intel® Xeon™ Processor 3.0Ghz with 800FSB 2MB Cache (Dual Processor Option)
- 4U Chassis with 16 hot-swap bays & 650W redundant power supply
- Intel® server board with Intel® E7520 (Lindenhurst) Chipset
- Kingston 1024MB DDR333 ECC Reg. (2x512MB)
- 1pc x Western Digital 250GB SATA RAID Drive
- 16 x 1" Hot-swap SATA drive bays
- Integrated ATI Rage XL SVGA PCI video controller
- Dual Intel® Gigabit Ethernet Controllers
- 2 SATA Channels with RAID 0 and 1

**\$3,499**

### 5U Advanced Storage Server SDR-5302T

The absolute best storage capacity available in a 5U format, at a incomparable cost/TB ratio



- Intel® Xeon™ Processor 3.0Ghz with 800FSB 2MB Cache (Dual Processor Option)
- 5U Chassis with 24 hot-swap bays & 950W redundant power supply
- Supermicro® server board with Intel® E7520 (Lindenhurst) Chipset
- Kingston 1024MB DDR2 400 ECC Reg. (2x512MB)
- 12pcs x Western Digital 250GB SATA RAID Drive
- 1pc x 3Ware 9550SX-12 port RAID Controller Card
- Integrated ATI Rage XL SVGA PCI video controller
- Dual Intel® 82541PI Single-Port Gigabit Ethernet Controllers

**\$5,999**

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

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

**1.877.727.7887 | WWW.SERVERSDIRECT.COM**

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





JON "MADDOG" HALL

# Shoring Up the Seawall

Trademark protection is necessary, if only because people seek to exploit what is unprotected.

**Hurricanes often show** the fragility of a beachhead town, particularly one whose seawall is not strong or high enough. Sometimes sandbags can save the day—if the weakness of the wall is local, if the strength of the hurricane is not high enough and if you have enough time and help to put the sandbags in place. But nothing beats having a well-designed and implemented seawall long before the storm hits, and often it is best to overbuild, as those “once in a lifetime” storms seem to be happening more and more often.

In 1994, a hurricane hit the Linux community. A man in Boston trademarked the term Linux and started to send out letters to Linux companies saying that he owned the term and would license it out to companies for *only* one-quarter of their revenue. The fledgling Linux community was shocked, and individual companies started to gather their sandbags (er, ah) lawyers to fight this attack. Linux International (LI) stepped in and acted as a channel for hiring the legal firm of Davis and Schroeder of Monterey, California (another beautiful coastal town). Through a long and arduous process costing tens of thousands of dollars and much pro bono (read that as gratis) legal service, LI



eventually had the trademark transferred to Linus Torvalds, who has held it ever since.

Linus wants everyone to use the name Linux for any legitimate purpose, and he really wanted it to go into the public domain. However, we found that there were people who wanted to use the name for a business that (although it was legal) was not what Linus wanted his name to be associated with—a porn site.

Linux had come into its own. As long as the word Linux had no value, no one cared about it. But as soon as the word Linux was perceived to be of value, people stepped forward to make what money they could, in both legitimate ways and less legitimate ways. We understand that other trademarks of other free and open-source software projects have seen similar issues, and we share their pain.

If the trademark had been in the public domain, there would have been nothing Linus could have done about the porn site, but because he was the registered owner of the mark, he could demand that the porn site stop using the mark, which they did after only a single, pointed letter from LI's law firm.

Eventually, the law firm convinced Linus that the trademark needed to be “protected” under the US trademark laws. This meant that if people just used the trademark any

way they wanted, eventually it would become public domain and all the pornography sites could be “Linuxporn”, and there was nothing he could do about it. Of course, there are other undesirable uses of the name, but this was one of the ones we encountered.

Because Linus considers himself an engineer, and did not want to deal with the day-to-day business of protecting the Linux trademark, he empowered the law firm and an old friend to create a nonprofit entity called Linux Mark Institute (LMI) to do this protection. LMI then started to sublicense the name to various companies who were using Linux as a trademark.

Although this column is too short to go into all the technical detail of what it means to use Linux as a trademark, or part of a trademark, the term Linux can be used in a fashion of fair use, which requires no licensing, but still should properly attribute the ownership of Linux to Linus Torvalds. I am sure the reader has seen the normal type of attribute at the bottom of some page (perhaps even in this magazine) that says, “The registered trademark Linux(R) is used pursuant to a license from Linus Torvalds, owner of the mark in the US and other countries.”

This is normally used when you just say the word Linux in some type of printed or electronic document, such as the *Linux Journal*.

But the incorporation of the name Linux into another name is what really needs to be licensed, and LMI's job became to seek out and sublicense the Linux mark to people and companies wishing to use the name.

LMI assembled an astute group of people to administer this trademark, people whose honesty and integrity were without question, and they started to formulate a sublicense that would:

- Protect the Linux mark.
- Allow businesses that were making money by using the name to help pay for the costs of administration.
- Allow LUGs, developers and low-revenue nonprofits to have a license at no charge.

And to make sure that things were on the right path, Linus kept the right to terminate this mechanism if it was not found to be carrying out his wishes.

There were some things that LMI absolutely needed to do to make sure it met the enforcement criteria. And, as often with free software things, the concept of licensing out the trademark for legitimate uses more or less did not mesh with the normal concept in trademark law of protecting the usage from other than the trademark holder “no matter

We understand that other trademarks of other free and open-source software projects have seen similar issues, and we share their pain.

what". LMI still managed to get through the knot holes to create a proper sublicense.

Just as everything started to move into place, an emergency happened in Australia that forced LMI to become active "before its time", and the world started to realize that although Linux (the kernel code) was already freely licensed, trademark law required the name to be licensed separately if not used in a fair-use way.

Immediately, there was wailing and gnashing of teeth from people who did not understand trademark law, did not make any use of the mark of Linux other than in fair use, and from people who never even heard of the word Linux before—in short, from everyone except those who were really affected. To be sure, some fine-tuning of the sublicense was needed, which was done, but not a single person who was really affected by the trademark sublicense objected to the premise, because as business people and trademark holders themselves, as members of the small commercial Linux community, they knew that the seawall needed rebuilding.

Today, people who use the term Linux as a proper trademark for their product can get a sublicense at the Linux Mark Institute ([www.linuxmark.org](http://www.linuxmark.org)). People and groups that make less than \$50,000 US per year in revenue pay nothing. Companies starting up that have not made any revenue to date using the Linux mark pay nothing. Only companies that have made more than \$50,000 US per year in revenue on products using the Linux mark pay a small percentage of that money to LMI. And, for very large companies making very large amounts of money on the Linux product, there is a cap on how much they pay.

Linux Mark Institute is a true nonprofit. None of the board of directors receives any salary. LMI employs only the bare minimum of staff. All money collected goes toward legal fees to protect the Linux name that otherwise would not be protected. Over time, if activity warrants it, we will reduce the license fees. But today, we feel the fees are fair, necessary and will not hurt anyone wishing to use the mark.

For all of those people who use the name Linux proudly, we ask only that you attribute it correctly, and that if you have any questions about whether you need a license, read the information at the Linux Mark Institute site.

Help us maintain a strong seawall. ■

---

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.

[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**®

Other company and product names are registered trademarks or trademarks of their respective owners. © 2006 FairCom Corporation



DOC SEARLS

# Use and Usefulness

Which comes first, the kernel chicken or the user-space egg?

I've always been intrigued by the distinctions between *kernel space* and *user space*. At the technical level, the distinction is largely between memory spaces: one where the kernel executes and provides services, and one where user processes run. As a rule, it's safer to run something in user space when it is possible, because user-space processes can't mess with the critical parts of the operating system. At a conceptual level, however, there also seems to be a distinction between *usefulness* and *use*.

I didn't start to see that distinction until I spent a week in October 2005 on a Linux Lunacy Geek Cruise with Andrew Morton, Ted Ts'o and a bunch of other kernel hackers. Andrew and Ted gave a number of talks, and I got a chance to spend additional time interviewing Andrew at length. As a result, I came to the conclusion that *Linux is a species*, and unpacked that metaphor in a Cruise Report on the *Linux Journal* Web site in November. Here's the gist of it:

Kernel development is not about Moore's Law. It's about natural selection, which is reactive, not proactive. Every patch to the kernel is adaptive, responding to changes in the environment as well as to internal imperatives toward general improvements on what the species is and does.

We might look at each patch, each new kernel version, even the smallest incremental ones, as a generation slightly better equipped for the world than its predecessors. Look at each patch submission—or each demand from a vendor that the kernel adapt to suit its needs in some way—as input from the environment to which the kernel might adapt.

We might look at the growth of Linux as that of a successful species that does a good job of adapting, thanks to a reproductive cycle that shames fruit flies. Operating systems, like other digital life forms, reproduce exuberantly. One `cp` command or `Ctrl-D`, and you've got a copy, ready to go—often into an environment where the species might be improved some more, patch by patch. As the population of the species grows and more patches come in, the kernel adapts and improves.

These adaptations are reactive more often than proactive. This is even, or perhaps especially, true for changes that large companies want. Companies such as IBM and HP, for example, might like to see proactive changes made to the kernel to better support their commercial applications.

Several years ago, I had a conversation with a Microsoft executive who told me that Linux had become a project of large commercial vendors, because so many kernel maintainers and contributors were employed by those vendors. Yet Andrew went out of his way to make clear, without irony, that the symbiosis between large vendors and the Linux kernel puts no commercial pressure on the kernel whatsoever. Each symbiote has its own responsibilities. To illustrate, he gave the case of one large com-

pany application: "The [application] team doesn't want to implement [something] until it's available in the kernel. One of the reasons I'd be reluctant to implement it in the kernel is that they haven't demonstrated that it's a significant benefit to serious applications. They haven't done the work to demonstrate that it will benefit applications. They're saying, 'We're not going to do the work if it's not in the kernel.' And I'm saying, 'I want to see that it will benefit the kernel if we put it in.'"

He added, "On the kernel team, we are concerned about the long-term viability and integrity of the code base. We're reluctant to put stuff in for specific reasons where a commercial company might do that." He says there is an "organic process" involved in vendor participation in the kernel.

It made my year when Greg Kroah-Hartman (a top-rank kernel maintainer) called this "one of the most insightful descriptions about what the Linux kernel really is, and how it is being changed over time".

A few weeks ago, I was talking with Don Marti about how all open-source projects seem to have the same kind of division between kernel space and user space—between code and dependencies on that code. It was in that conversation that I realized the main distinction was between usefulness and use. Roles as well as purposes were involved. Only developers contribute code. The influence of users, or even "usability experts", is minimized by the meritocracy that comprises the development team. "Show me the code" is a powerful filter.

Most imperatives of commercial development originate and live in user space. These include selling products, making profits and adding product features that drive future sales. None of these motivations are of much (if any) interest to kernel development. Again, kernel development is reactive, not proactive. For companies building on Linux, the job is putting Linux to use, not telling it how to be useful. Unless, of course, you have useful code to contribute. (Greg Kroah-Hartman has put together an excellent set of recommendations. See the on-line Resources for links.)

A few tradeshows ago, Dan Frye of IBM told me it took years for IBM to discover that the company needed to adapt to its kernel developers, rather than vice versa. I am sure other employers of kernel developers have been making the same adjustments. How long before the rest of the world follows? And what will the world learn from that adjustment that it doesn't know now?

I began to see an answer take shape at O'Reilly's Emerging Technology Conference in March 2006. I was sitting in the audience, writing and rewriting this very essay, when George Dyson took the stage and blew my mind. George grew up in Princeton, hanging around the Institute for Advanced Study where his father, Freeman Dyson, worked with Godel, Einstein, Turing, von Neumann and other legends in mathematics, physics and computing. Today, George is a historian studying the work of those same great minds, plus antecedents running back hundreds of years.

George's lecture, titled "Turing's Cathedral", reviews the deep history of computing, its supportive mathematics and the staging

# SUPERMICRO®

*Faster · Higher Capacity · Future Trend*

## SAS over SCSI

*Complete with SES2 & RAID5 Solutions*



### 1U SuperServer 6014P-32R

- ✓ Dual Intel® Dual-core 64-bit Xeon™ processors support
- ✓ SUPER® X6DHP-3G2 Serverboard
- ✓ Intel® E7520 chipset/800MHz
- ✓ Up to 16GB ECC Reg. DDRII 400 SDRAM
- ✓ 1 Universal PCI-X 133MHz (full-length)/1 PCI-Express x8
- ✓ 1 Universal PCI-X 133MHz (low-profile)/1 PCI-Express x8
- ✓ Dual Gigabit LAN ports
- ✓ Adaptec 9410W controller for 8 SAS/SATA drives
- ✓ AOC-LPZCR2 (Zero-Channel RAID) support
- ✓ 4 x 1" hot-swap SAS/SATA drive bays w/SES2
- ✓ 1 slim floppy & 1 slim DVD-ROM drive
- ✓ 600W cold-swap power supply w/I<sup>2</sup>C
- ✓ 5x 4cm heavy duty fans w/optimal fan speed control
- ✓ System management: Super® Doctor III & IPMI 2.0

\* RAID 0, 1, 5, 10 support

### 2U SuperServer 6024H-32R

- ✓ Dual Intel® Dual-core 64 bit Xeon™ processors support
- ✓ SUPER® X6DH3-G2 Serverboard
- ✓ Intel® E7520 chipset/800MHz
- ✓ Up to 16GB ECC Reg. DDRII 400 SDRAM
- ✓ 1/2 64-bit PCI-X 133/100MHz (full-size)
- ✓ Dual Gigabit LAN ports
- ✓ Adaptec 9410W controller for 8 SAS/SATA drives
- ✓ AOC-LPZCR2 (Zero-Channel RAID) support
- ✓ 6 x 1" hot-swap SAS/SATA drive bays w/SES2
- ✓ 1 x 5.25" drive bay, 1 floppy drive & 1 slim CD-ROM
- ✓ 500W redundant power supply w/I<sup>2</sup>C
- ✓ 4 x 8cm heavy duty fans w/optimal fan speed control & air shroud
- ✓ System management: Super® Doctor III & IPMI 2.0

### 4U SuperServer 7044H-32R

- ✓ Dual Intel® Dual-core 64 bit Xeon™ processors support
- ✓ SUPER® X6DH3-G2 Serverboard
- ✓ Intel® E7520 chipset/800MHz
- ✓ Up to 16GB ECC Reg. DDRII 400 SDRAM
- ✓ 2 PCI Express x8
- ✓ 1/2 64-bit PCI-X 133/100MHz & 1 PCI
- ✓ Dual Gigabit LAN ports
- ✓ Adaptec 9410W controller for 8 SAS/SATA drives
- ✓ AOC-LPZCR2 (Zero-Channel RAID) support
- ✓ 8 x 1" hot-swap SAS/SATA drive bays w/SES2
- ✓ 90° rotatable module: 2 USB ports, FDD, 2 x 5.25" drive bays
- ✓ 760W triple-redundant power supply w/I<sup>2</sup>C
- ✓ 100% cooling redundancy\*\*: 4 x 8cm, 2 x 8cm fans w/optimal fan speed control & air shroud
- ✓ System management: Super® Doctor III & IPMI 2.0

\*\* System remains at 100% performance even if a key fan fails.

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

Arrow Electronics  
1-888-427-2250  
www.arrowna.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

of a shift in computing from the mechanical to the biological—one that von Neumann had begun to expect when he died tragically in 1957 at the age of 53. Here's how George approaches questions similar to the one that had been on my mind:

"The whole human memory can be, and probably in a short time will be, made accessible to every individual", wrote H. G. Wells in his 1938 prophecy *World Brain*. "This new all-human cerebrum need not be concentrated in any one single place. It can be reproduced exactly and fully, in Peru, China, Iceland, Central Africa, or wherever else seems to afford an insurance against danger and interruption. It can have at once, the concentration of a craniate animal and the diffused vitality of an amoeba." Wells foresaw not only the distributed intelligence of the World Wide Web, but the inevitability that this intelligence would coalesce, and that power, as well as knowledge, would fall under its domain. "In a universal organization and clarification of knowledge and ideas...in the evocation, that is, of what I have here called a World Brain...in that and in that alone, it is maintained, is there any clear hope of a really Competent Receiver for world affairs....We do not want dictators, we do not want oligarchic parties or class rule, we want a widespread world intelligence conscious of itself."

## It made my year when Greg Kroah-Hartman (a top-rank kernel maintainer) called this "one of the most insightful descriptions about what the Linux kernel really is, and how it is being changed over time".

Then:

In the early 1950s, when mean time between memory failure was measured in minutes, no one imagined that a system depending on every bit being in exactly the right place at exactly the right time could be scaled up by a factor of  $10^{13}$  in size, and down by a factor of  $10^6$  in time. Von Neumann, who died prematurely in 1957, became increasingly interested in understanding how biology has managed (and how technology might manage) to construct reliable organisms out of unreliable parts. He believed the von Neumann architecture would soon be replaced by something else. Even if codes could be completely debugged, million-cell memories could never be counted upon, digitally, to behave consistently from one kilocycle to the next.

Fifty years later, thanks to solid state micro-electronics, the von Neumann matrix is going strong. The problem has shifted from how to achieve reliable results using sloppy hardware, to how to achieve reliable results using sloppy code. The von Neumann architecture is here to stay. But new forms of architecture, built upon the underlying layers of Turing-von Neumann machines, are starting to grow. What's next? Where was von Neumann heading when his program came to a halt?

This is all excerpted from an earlier lecture by George, by the same title as the one he was giving at eTech. In this earlier lecture, George was focused on AI:

I found myself recollecting the words of Alan Turing, in his seminal paper "Computing Machinery and Intelligence", a founding document in the quest for true AI. "In attempt-

ing to construct such machines we should not be irreverently usurping His power of creating souls, any more than we are in the procreation of children", Turing had advised. "Rather we are, in either case, instruments of His will providing mansions for the souls that He creates."

Then he added, "Google is Turing's cathedral, awaiting its soul."

I think, however, the cathedral is bigger than Google. In fact, I think it's bigger than a cathedral. I think it's a new world, built on materials no less natural yet man-made than the rocks and wood shaped and assembled into nave and transept, buttress and spire.

I reached that conclusion watching George flash quote after quote up on the screen in the front of the ballroom, each drawing another line in the shape we came to call computing. I photographed as many as I could, and transcribed a number of them. I've arranged them in chronological order, starting 450 years ago, with several more added in from my own quote collection. Follow the threads:

- "Why may we not say that all Automata (Engines that move themselves by springs and wheeles as doth a watch) have artificiall life?"—Thomas Hobbes, 1651
- "By Ratiocination, I mean computation. Now to compute, is either to collect the sum of many things that are added together, and to know what remains when one thing is taken out of another...and if any man adde Multiplication and Division, I will not be against it, seeing Multiplication is nothing but Addition of equals one to another, and Division is nothing but a Subtraction of equals one from another, as often as is possible. So that all Ratiocination is comprehended in these two operations of the minde Addition and Subtraction."—Thomas Hobbes, 1656
- "This [binary] calculus could be implemented by a machine (without wheels)...provided with holes in such a way that they can be opened and closed. They are to be open at those places that correspond to a 1 and remain closed at those that correspond to a 0. Through the opened gates small cubes or marbles are to fall into tracks, through the others nothing. It [the gate array] is to be shifted from column to column as required."—G.W. von Leibniz, March 16, 1679
- "Is it a fact—or have I dreamed it—that, by means of electricity, the world of matter has become a great nerve, vibrating thousands of miles in a breathless point of time? Rather the round globe is a vast head, a brain, instinct with intelligence! Or shall I say, it is itself a thought, nothing but a thought, and no longer the substance which we deemed it?"—Nathaniel Hawthorne, 1851
- "I see the Net as a world we might see as a bubble. A sphere. It's growing larger and larger, and yet inside, every point in that sphere is visible to every other one. That's the architecture of a sphere. Nothing stands between any two points. That's its virtue: it's empty in the middle. The dis-

tance between any two points is functionally zero, and not just because they can see each other, but because nothing interferes with operation between any two points. There's a word I like for what's going on here: terraform. It's the verb for creating a world. That's what we're making here: a new world. Now the question is, what are we going to do to cause planetary existence? How can we terraform this new world in a way that works for the world and not just ourselves?"—Craig Burton, in *Linux Journal*, 1999

■ "Here are three basic rules of behavior that are tied directly to the factual nature of the Internet: 1) No one owns it. 2) Everyone can use it. 3) Anyone can improve it."—"World of Ends", by Doc Searls and David Weinberger, 2003

■ "There are a couple of reasons why we have national parks and access to the seashore. Some things are so much the gifts of nature that they should be reserved for everyone. And some things (like the sea, and like the Internet) are so important to each of us that keeping them freely available makes us a group of citizens rather than slaves...Now—the Internet wasn't created by nature; it's an agreement between machines made possible by the designers of that agreement (or protocol). But it is a great gift, and it is very important to being a citizen, and for these reasons it is owned by all for common use. It's a commons, like the Boston Common. And no sovereign ever showed up to which the people who 'own' the Internet (that is, everyone) surrendered their ownership."—Susan Crawford, January 2003

■ "We had this idea back in the 70s that one day we would make computers that would somehow be intelligent on their own. And it's not quite working that way. What we're doing is making computers intelligent because we're part of them."—Tim O'Reilly at eTech 2006

As creatures, human beings are gifted with something perhaps even more significant than the powers of intelligence and speech. We also have the capacity to extend the boundaries of our bodies beyond our skin, hair and nails. Through a process of indwelling, we are enlarged and empowered by our clothes, tools and vehicles. When we grab a hammer and drive a nail, the hammer becomes an extension of our arm. Our senses extend through the wood of the handle and the metal of the head, as we pound a nail through a board. Oddly, the hammer does not make us superhuman, but more human. Because nothing could be more human than to use a tool.

Likewise, when we drive a car, ride a bike or pilot a plane, our senses extend to mechanical perimeters. We don't just think "my tires", "my wings", "my fender", "my engine". We know these things are ours. They are parts of our selves, enlarged by the merging of sense and skill and material.

A robin is born knowing how to build a nest. A human is born knowing how to do little beyond suckling. Yet because we are gifted with an endless capacity for learning, and for enlarging our selves, and for doing these things together in groups of all sizes, we have built something larger than ourselves called civilization.

Open-source infrastructural building materials and methods have enabled us to build a new framework, a new environment, for civilization. Call it a giant brain, a World of Ends, or a network of networks. In every case, it is a product of the form of nature we call human.

The purpose of this new world—this natural environment for business, study, games and countless other human activities—is to be useful. In the same way that our senses extend from our bodies to our tools and vehicles, the usefulness of kernel-space code extends into the Net that's built on that code.

As a result, user space has become almost unimaginably large. And sure to become larger. ■

**Resources for this article:** [www.linuxjournal.com/article/8942](http://www.linuxjournal.com/article/8942).

Doc Searls is Senior Editor of *Linux Journal*.



"Company With A Vision"

1-877-25-SERVER : [www.genstor.com](http://www.genstor.com)

**GREAT MINDS – GREAT SOLUTIONS**  
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

## STORAGE

**SATA-NAS-DAS-iSCSI-SAS Storage Solutions**



5U Dual Xeon/  
Opteron SATA Storage  
Upto 24GB RAM  
Upto 26 SATA HDD  
Upto 13TB Storage  
Starting @ \$ 3930.00  
Call for your custom needs

**CALL FOR BEST PRICED iSCSI SOLUTIONS**

## CUSTOM APPLIANCE SOLUTION

**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

**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.

## Thinstuff's Touch Rdpserver

The Slovakian company Thinstuff recently released Touch Rdpserver, a Linux-based Remote Desktop Protocol (RDP) terminal server solution for RDP 5.0, 5.1, 5.2, Windows CE and rdesktop clients. Thinstuff claims that the product provides virtual X servers for highly optimized X11 to RDP translation, a server management framework for operating large-scale terminal server clusters and low-bandwidth consumption. Thinstuff offers a free, downloadable demo on its Web site.

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



## SlickEdit v11

If your mission is to code in different operating systems, SlickEdit v11 might be your editor. This latest update, from its eponymous mother company, adds a range of new features, such as code templates, comment wrapping, comment autogeneration for Javadoc and XMLdoc, regex evaluator, Xft font support, Vim emulation, universal binary support for Macintosh and more. These build on legacy features such as the Context Tagging engine and the DIFFzilla tool for comparing files and directories. SlickEdit supports more than 40 languages on seven platforms, including Linux kernel 2.4 and later. A free trial edition is available for download on the company's Web site.

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

## Airchitex's Cuckoo

Airchitex's Cuckoo is a network appliance for those who want the peace of mind that comes with having one's own SNTP time server. Cuckoo is designed to supply accurate and precise time to every machine in your network, sans loud bird. The device receives its time information from global positioning system (GPS) satellites, making it accurate to +/- 1 millisecond. Network protocols include NTP, TIME and DAYTIME; autosecs for time zones and daylight-savings time work in the US. Airchitex claims that Cuckoo avoids common time errors stemming from PC clock drift and Internet routing latency, among other issues.


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



## Centrify's DirectControl Suite 3.0

Grudgingly or not, many of us use Microsoft's Active Directory for user authentication and system access. Against this backdrop, we are gleefully adding desktops of various flavors. The detriment to this variety comes in the form of cross-platform identity and access management issues. Centrify's DirectControl Suite attempts to eliminate these issues by delivering secure access control and centralized identity management by seamlessly integrating UNIX, Linux, Mac OS, J2EE and Web platforms with Microsoft Active Directory. Admins can control who has access to what systems and applications while centrally managing a consistent, global security policy. End users can access all systems with a single password and can get access in minutes rather than days.

[www.centrify.com](http://www.centrify.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.



## Want to increase your ability to administer large numbers of PCs with a minimum amount of people?

### Here's why.

Traditional Linux server-client architectures—especially thin client—are complicated, cumbersome approaches.

They require lots of servers (about one for every 30 clients), constant maintenance, teams of administrators and large capital investments.

They are also expensive, inefficient, complicated and not terribly secure.

#### **IDEAS<sup>sm</sup>. The better way**

Integrity Diskless Enterprise Architecture Solutions (IDEAS)—based on Red Hat Enterprise Server V3 and 4-32 and 64-bit OS—is the simpler, more effective way to build diskless desktop computing.

It streamlines installation, scalability and system administration.

A single IDEAS OS server easily handles hundreds of users. Fewer servers means fewer system administrators.

Yet systems remain extremely powerful, easily handling the most demanding development and visualization applications without affecting operations of other OS-based systems on the network.

1-800-657-9352 [www.IntegrityLinux.com](http://www.IntegrityLinux.com)

## Then you need IDEAS<sup>sm</sup>

### Integrity Diskless Enterprise Architecture Solutions

#### **CREM security**

Diskless computing is an ideal solution for security-intensive operations because it eliminates classified removable electronic media (CREM). It's the ultimate security because there's no hard drive to track or manage.

#### **A simple technology whose time has come**

Diskless computing simplifies network computing (including implementation, maintenance, security and administration) without managing disk images. As a result, it improves your total cost of ownership (TCO).

To find out how diskless computing can simplify your network, come see us at [www.IntegrityLinux.com](http://www.IntegrityLinux.com)...

**and discover a new world of IDEAS<sup>sm</sup>.**

**INTEGRITY**  
Linux Systems Division



RAID and Logical Volume Managers are great,

until you lose data.

RICHARD BULLINGTON-MCGUIRE

Photo by Mehdi Pasha Oskooie

# Recovery of RAID and LVM2 Volumes

**The combination of** Linux software RAID (Redundant Array of Inexpensive Disks) and LVM2 (Logical Volume Manager, version 2) offered in modern Linux operating systems offers both robustness and flexibility, but at the cost of complexity should you ever need to recover data from a drive formatted with software RAID and LVM2 partitions. I found this out the hard way when I recently tried to mount a system disk created with RAID and LVM2 on a different computer. The first attempts to read the filesystems on the disk failed in a frustrating manner.

I had attempted to put two hard disks into a small-form-factor computer that was really designed to hold only one hard disk, running the disks as a

mirrored RAID 1 volume. (I refer to that system as raidbox for the remainder of this article.) This attempt did not work, alas. After running for a few hours, it would power-off with an automatic thermal shutdown failure. I already had taken the system apart and started re-installing with only one disk when I realized there were some files on the old RAID volume that I wanted to retrieve.

Recovering the data would have been easy if the system did not use RAID or LVM2. The steps would have been to connect the old drive to another computer, mount the filesystem and copy the files from the failed volume. I first attempted to do so, using a computer I refer to as recoverybox, but this attempt met with frustration.

# When “Medium” just doesn’t cut it.



**Q14 - Dual Socket Dual Core AMD Opteron - 4 Cores**  
AMD64 Opteron Dual-Core 64-bit Technology  
or 64-bit Intel® Xeon® processors  
Up to 32GB ECC Registered Memory  
Up to 2TB of RAID Storage with 4 Hot Swap Drives

**L**




**Q28 - Dual Socket Dual Core AMD Opteron - 4 Cores**  
AMD64 Opteron Dual-Core Technology  
or 64-bit Intel® Xeon® processors  
Up to 32GB 400MHz ECC Memory  
Up to 4TB of RAID Storage with 8 Hot Swap Drives

**XL**



**Q416 - Quad Socket / Dual Core AMD Opteron - 8 Cores**  
AMD64 Opteron Dual-Core Technology for up to 8 Cores  
or 64-bit Intel® Xeon® processors  
Up to 128GB ECC Registered Memory  
Up to 8TB of RAID Storage with 16 Hot Swap Drives  
XXXL also available with Q524 and 24 Drives for up to 12TB RAID

**XXL**



**QSOL.COM**  
Server Appliances  
**800.933.7510**  
*Serving Servers since 1999*

## Why Was This So Hard?

Getting to the data proved challenging, both because the data was on a logical volume hidden inside a RAID device, and because the volume group on the RAID device had the same name as the volume group on the recovery system.

Some popular modern operating systems (for example, Red Hat Enterprise Linux 4, CentOS 4 and Fedora Core 4) can partition the disk automatically at install time, setting up the partitions using LVM for the root device. Generally, they set up a volume group called VolGroup00, with two logical volumes, LogVol00 and LogVol01, the first for the root directory and the second for swap, as shown in Listing 1.

The original configuration for the software RAID device had three RAID 1 devices: md0, md1 and md2, for /boot, swap and /, respectively. The LVM2 volume group was on the biggest RAID device, md2. The volume group was named VolGroup00. This seemed like a good idea at the time, because it meant that the partitioning configuration for this box looked similar to how the distribution does things by default. Listing 2 shows how the software RAID array looked while it was operational.

If you ever name two volume groups the same thing, and something goes wrong, you may be faced with the same problem. Creating conflicting names is easy to do, unfortunately, as the operating system has a default primary volume group name of VolGroup00.

## Restoring Access to the RAID Array Members

To recover, the first thing to do is to move the drive to another machine. You can do this pretty easily by putting the drive in a USB2 hard drive enclosure. It then will show up as a SCSI hard disk device, for example, /dev/sda, when you plug it in to your recovery computer. This reduces the risk of damaging the recovery machine while attempting to install the hardware from the original computer.

The challenge then is to get the RAID setup recognized and to gain access to the logical volumes within. You can use `sfdisk -l /dev/sda` to check that the partitions on the old drive are still there.

To get the RAID setup recognized, use `mdadm` to scan the devices for their raid volume UUID signatures, as shown in Listing 3.

This format is very close to the format of the `/etc/mdadm.conf` file that the `mdadm` tool uses. You need to redirect the output of `mdadm` to a file, join the device lines onto the ARRAY lines and put in a nonexistent second device to get a RAID1 configuration. Viewing the md array in degraded mode will allow data recovery:

```
[root@recoverybox ~]# mdadm --examine --scan /dev/sda1
↳ /dev/sda2 /dev/sda3 >> /etc/mdadm.conf
[root@recoverybox ~]# vi /etc/mdadm.conf
```

Edit `/etc/mdadm.conf` so that the devices statements are on the same lines as the ARRAY statements, as they are in Listing 4. Add the "missing" device to the devices entry for each array member to fill out the raid1 complement of two devices per array. Don't forget to renumber the md entries if the recovery computer already has md devices and ARRAY statements in `/etc/mdadm.conf`.

Then, activate the new md devices with `mdadm -A -s`, and check `/proc/mdstat` to verify that the RAID array is active. Listing 5 shows how the RAID array should look.

If md devices show up in `/proc/mdstat`, all is well, and you can move on to getting the LVM volumes mounted again.

### Listing 1.

## TYPICAL LVM DISK CONFIGURATION

```
[root@recoverybox ~]# /sbin/sfdisk -l /dev/hda

Disk /dev/hda: 39560 cylinders, 16 heads, 63 sectors/track
Warning: The partition table looks like it was made
         for C/H/S=*/255/63 (instead of 39560/16/63).
For this listing I'll assume that geometry.
Units = cylinders of 8225280 bytes, blocks of 1024 bytes, counting from 0

   Device Boot  Start      End  #cyls   #blocks  Id System
/dev/hda1    *      0+       12     13-     104391   83 Linux
/dev/hda2                13     2481    2469    19832242+  8e Linux LVM
/dev/hda3                0         -         0         0     0 Empty
/dev/hda4                0         -         0         0     0 Empty
[root@recoverybox ~]# /sbin/pvscan
PV /dev/hda2   VG VolGroup00   lvm2 [18.91 GB / 32.00 MB free]
Total: 1 [18.91 GB] / in use: 1 [18.91 GB] / in no VG: 0 [0 ]
[root@recoverybox ~]# /usr/sbin/lvscan
ACTIVE        '/dev/VolGroup00/LogVol00' [18.38 GB] inherit
ACTIVE        '/dev/VolGroup00/LogVol01' [512.00 MB] inherit
```

### Listing 2.

## SOFTWARE RAID DISK CONFIGURATION

```
[root@raidbox ~]# /sbin/sfdisk -l /dev/hda

Disk /dev/hda: 9729 cylinders, 255 heads, 63 sectors/track
Units = cylinders of 8225280 bytes, blocks of 1024 bytes, counting from 0

   Device Boot  Start      End  #cyls   #blocks  Id System
/dev/hda1    *      0+       12     13-     104391   fd Linux raid
autodetect
/dev/hda2                13       77      65     522112+   fd Linux raid
autodetect
/dev/hda3                78     9728    9651    77521657+ fd Linux raid
autodetect
/dev/hda4                0         -         0         0     0 Empty
[root@raidbox ~]# cat /proc/mdstat
Personalities : [raid1]
md2 : active raid1 hdc3[1] hda3[1]
      77521536 blocks [2/2] [UU]

md1 : active raid1 hdc2[1] hda2[1]
      522048 blocks [2/2] [UU]

md0 : active raid1 hdc1[1] hda1[1]
      104320 blocks [2/2] [UU]
```

### Listing 3.

## SCANNING A DISK FOR RAID ARRAY MEMBERS

```
[root@recoverybox ~]# mdadm --examine --scan /dev/sda1 /dev/sda2 /dev/sda3
ARRAY /dev/md2 level=raid1 num-devices=2
↳ UUID=532502de:90e44fb0:242f485f:f02a2565
  devices=/dev/sda3
ARRAY /dev/md1 level=raid1 num-devices=2
↳ UUID=75fa22aa:9a11bcad:b42ed14a:b5f8da3c
  devices=/dev/sda2
ARRAY /dev/md0 level=raid1 num-devices=2
↳ UUID=b3cd99e7:d02be486:b0ea429a:e18ccf65
  devices=/dev/sda1
```

Listing 4.

## /etc/mdadm.conf

```
DEVICE partitions
ARRAY /dev/md0 level=raid1 num-devices=2
  ↳UUID=b3cd99e7:d02be486:b0ea429a:e18ccf65
  ↳devices=/dev/sda1,missing
ARRAY /dev/md1 level=raid1 num-devices=2
  ↳UUID=75fa22aa:9a11bcad:b42ed14a:b5f8da3c
  ↳devices=/dev/sda2,missing
ARRAY /dev/md2 level=raid1 num-devices=2
  ↳UUID=532502de:90e44fb0:242f485f:f02a2565
  ↳devices=/dev/sda3,missing
```

Listing 5.

## Reactivating the RAID Array

```
[root@recoverybox ~]# mdadm -A -s
[root@recoverybox ~]# cat /proc/mdstat
Personalities : [raid1]
md2 : active raid1 sda3[1]
      77521536 blocks [2/1] [_U]

md1 : active raid1 sda2[1]
      522048 blocks [2/1] [_U]

md0 : active raid1 sda1[1]
      104320 blocks [2/1] [_U]

unused devices: <none>
```

## Recovering and Renaming the LVM2 Volume

The next hurdle is that the system now will have two sets of lvm2 disks with VolGroup00 in them. Typically, the `vgchange -a -y` command would allow LVM2 to recognize a new volume group. That won't work if devices containing identical volume group names are present, though. Issuing `vgchange -a -y` will report that VolGroup00 is inconsistent, and the VolGroup00 on the RAID device will be invisible. To fix this, you need to rename the volume group that you are about to mount on the system by hand-editing its lvm configuration file.

If you made a backup of the files in /etc on raidbox, you can edit a copy of the file /etc/lvm/backup/VolGroup00, so that it reads VolGroup01 or RestoreVG or whatever you want it to be named on the system you are going to restore under, making sure to edit the file itself to rename the volume group in the file.

If you don't have a backup, you can re-create the equivalent of an LVM2 backup file by examining the LVM2 header on the disk and editing out the binary stuff. LVM2 typically keeps copies of the metadata configuration at the beginning of the disk, in the first 255 sectors following the partition table in sector 1 of the disk. See /etc/lvm/lvm.conf and man lvm.conf for more details. Because each disk sector is typically 512 bytes, reading this area will yield a 128KB file. LVM2 may have stored several different text representations of the LVM2 configuration stored on the partition itself in the first 128KB. Extract these to an ordinary file as follows, then edit the file:

```
dd if=/dev/md2 bs=512 count=255 skip=1 of=/tmp/md2-raw-start
vi /tmp/md2-raw-start
```

You will see some binary gibberish, but look for the bits of plain text. LVM treats this metadata area as a ring buffer, so there may be multiple configuration entries on the disk. On my disk, the first entry had only the details for the physical volume and volume group, and the next entry had the logical volume information. Look for the block



# ASA COMPUTERS

## Want your business to be more productive?

The ASA Servers powered by the Intel® Xeon™ Processor provides the quality and dependability to keep up with your growing business.

## Hardware Systems For The Open Source Community—Since 1989

(Linux, FreeBSD, NetBSD, OpenBSD, Solaris, MS, etc.)



### 6TB + in 5U—\$7,599

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

### 1U Dual Itanium IDE—\$3,615

Dual Intel® Itanium® 2 1.4 Ghz  
2 GB ECC DDR  
1 of 4 x 40 GB HDD  
Dual Gigabit LAN  
Based on Supermicro 6113M-i



### 14" Deep Appliance Server—\$865

Intel® Xeon™ 2.4 Ghz Processor  
40 GB Hard Drive, One GigE  
Options: CD, FD, 2nd HD, Your Logo  
on Bezel  
**Call for Low Cost Options.**

### 1U Dual Xeon™ EM64T Superserver—\$1,605

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



### Your Custom Appliance Solution

Let us know your needs, we will get you a solution

"Your Logo Here"



### ASA Collocation

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

### ASA Collocation Special

First month of collocation free.\*

### Storage Solutions

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

### Clusters

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

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



2354 Calle Del Mundo,  
Santa Clara, CA 95054

www.asacomputers.com

Email: sales@asacomputers.com

P: 1-800-REAL-PCS | FAX: 408-654-2910



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

Prices and availability subject to change without notice. Not responsible for typographical errors.

#### Listing 6.

### Modified Volume Group Configuration File

```
VolGroup01 {
id = "xQZqTG-V4wn-DLeQ-bJ0J-GEHB-4teF-A4PPBv"
seqno = 1
status = ["RESIZEABLE", "READ", "WRITE"]
extent_size = 65536
max_lv = 0
max_pv = 0

physical_volumes {

pv0 {
id = "tRACEy-cstP-kk18-zQFZ-ErG5-QAIV-YqHItA"
device = "/dev/md2"

status = ["ALLOCATABLE"]
pe_start = 384
pe_count = 2365
}
}

# Generated by LVM2: Sun Feb  5 22:57:19 2006
```

#### Listing 7.

### Activating the Recovered LVM2 Volume

```
[root@recoverybox ~]# vgcfgrestore -f VolGroup01 VolGroup01
[root@recoverybox ~]# vgscan
  Reading all physical volumes.  This may take a while...
  Found volume group "VolGroup01" using metadata type lvm2
  Found volume group "VolGroup00" using metadata type lvm2
[root@recoverybox ~]# pvscan
PV /dev/md2      VG VolGroup01   lvm2 [73.91 GB / 32.00 MB free]
PV /dev/hda2    VG VolGroup00   lvm2 [18.91 GB / 32.00 MB free]
Total: 2 [92.81 GB] / in use: 2 [92.81 GB] / in no VG: 0 [0 ]
[root@recoverybox ~]# vgchange VolGroup01 -a y
  1 logical volume(s) in volume group "VolGroup01" now active
[root@recoverybox ~]# lvscan
ACTIVE            '/dev/VolGroup01/LogVol00' [73.88 GB] inherit
ACTIVE            '/dev/VolGroup00/LogVol00' [18.38 GB] inherit
ACTIVE            '/dev/VolGroup00/LogVol01' [512.00 MB] inherit
```

of text with the most recent timestamp, and edit out everything except the block of plain text that contains LVM declarations. This has the volume group declarations that include logical volumes information. Fix up physical device declarations if needed. If in doubt, look at the existing `/etc/lvm/backup/VolGroup00` file to see what is there. On disk, the text entries are not as nicely formatted and are in a different order than in the normal backup file, but they will do. Save the trimmed configuration as `VolGroup01`. This file should then look like Listing 6.

Once you have a volume group configuration file, migrate the volume group to this system with `vgcfgrestore`, as Listing 7 shows.

At this point, you can now mount the old volume on the new system, and gain access to the files within, as shown in Listing 8.

Now that you have access to your data, a prudent final step would be to back up the volume group information with `vcfgbackup`, as Listing 9 shows.

#### Listing 8.

### Mounting the Recovered Volume

```
[root@recoverybox ~]# mount /dev/VolGroup01/LogVol00 /mnt
[root@recoverybox ~]# df -h
Filesystem                Size  Used Avail Use% Mounted on
/dev/mapper/VolGroup00-LogVol00
                          19G  4.7G   13G   28% /
/dev/hda1                  99M   12M   82M   13% /boot
none                       126M    0   126M    0% /dev/shm
/dev/mapper/VolGroup01-LogVol00
                          73G  2.5G   67G    4% /mnt

# ls -l /mnt
total 200
drwxr-xr-x  2 root root  4096 Feb  6 02:36 bin
drwxr-xr-x  2 root root  4096 Feb  5 18:03 boot
drwxr-xr-x  4 root root  4096 Feb  5 18:03 dev
drwxr-xr-x 79 root root 12288 Feb  6 23:54 etc
drwxr-xr-x  3 root root  4096 Feb  6 01:11 home
drwxr-xr-x  2 root root  4096 Feb  21  2005 initrd
drwxr-xr-x 11 root root  4096 Feb  6 02:36 lib
drwx-----  2 root root 16384 Feb  5 17:59 lost+found
drwxr-xr-x  3 root root  4096 Feb  6 22:12 media
drwxr-xr-x  2 root root  4096 Oct  7 09:03 misc
drwxr-xr-x  2 root root  4096 Feb  21  2005 mnt
drwxr-xr-x  2 root root  4096 Feb  21  2005 opt
drwxr-xr-x  2 root root  4096 Feb  5 18:03 proc
drwxr-xr-x  5 root root  4096 Feb  7 00:19 root
drwxr-xr-x  2 root root 12288 Feb  6 22:37 sbin
drwxr-xr-x  2 root root  4096 Feb  5 23:04 selinux
drwxr-xr-x  2 root root  4096 Feb  21  2005 srv
drwxr-xr-x  2 root root  4096 Feb  5 18:03 sys
drwxr-xr-x  3 root root  4096 Feb  6 00:22 tftpboot
drwxrwxrwt  5 root root  4096 Feb  7 00:21 tmp
drwxr-xr-x 15 root root  4096 Feb  6 22:33 usr
drwxr-xr-x 20 root root  4096 Feb  5 23:15 var
```

#### Listing 9.

### Backing Up Recovered Volume Group Configuration

```
[root@teapot-new ~]# vgcfgbackup
Volume group "VolGroup01" successfully backed up.
Volume group "VolGroup00" successfully backed up.
[root@teapot-new ~]# ls -l /etc/lvm/backup/
total 24
-rw-----  1 root root 1350 Feb 10 09:09 VolGroup00
-rw-----  1 root root 1051 Feb 10 09:09 VolGroup01
```

## Conclusion

LVM2 and Linux software RAID make it possible to create economical, reliable storage solutions with commodity hardware. One trade-off involved is that some procedures for recovering from failure situations may not be clear. A tool that reliably extracted old volume group information directly from the disk would make recovery easier. Fortunately, the designers of the LVM2 system had the wisdom to keep plain-text backup copies of the configuration on the disk itself. With a little patience and some research, I was able to regain access to the logical volume I thought was lost; may you have as much success with your LVM2 and RAID installation. ■

**Resources for this article:** [www.linuxjournal.com/article/8948](http://www.linuxjournal.com/article/8948).

Richard Bullington-McGuire is the Managing Partner of PKR Internet, LLC, a software and systems consulting firm in Arlington, Virginia, specializing in Linux, Open Source and Java. He has been a Linux sysadmin since 1994. You can reach him at [rbullington@pkrinternet.com](mailto:rbullington@pkrinternet.com).

**Cool. Efficient. Reliable.**



**Meet your  
data center dream date.**



With challenges ranging from power shortages to thermal issues, today's data center managers are seeking cooler, more efficient and more reliable solutions.

Rackable Systems' award-winning servers and storage products are designed to reduce power consumption and increase reliability. Customers deploying Rackable Systems servers with DC power and the AMD Opteron™ processor experience monthly power savings of up to 30% and heat reduction in the chassis of up to 40%. And with density levels as high as 88 dual-core servers per cabinet, this all adds up to the most data center-friendly solution out there.

Time to meet your dream date.

To learn more and set up a first date, visit [www.rackable.com/dreamdate](http://www.rackable.com/dreamdate).



408.240.8300 [www.rackable.com](http://www.rackable.com)

# Network Transparency

KIO makes it easy to access network storage from Konqueror and many other KDE KIO-enabled applications.

with **KIO**

Jes Hall

Kioslaves are out-of-process protocol handling plugins that let you access various services as though they were part of your local filesystem. They can be accessed through the Konqueror file manager and any standard file dialog in most KDE applications. The applications using Kio don't need to be aware of how to access the remote host or device—the ioslave handles it all. This provides powerful and flexible network transparency for KDE applications.

KDE includes a large range of basic kioslaves. Some KDE add-on applications install their own. You can add new kioslaves at any time, and all KDE KIO-aware applications will work with them. If you use Konqueror as a Web browser, you're already using KIO, as the HTTP, HTTPS and FTP protocols are implemented as kioslaves.

## Fish

The KDE fish kioslave provides a graphical interface for manipulating files over Secure SHell (SSH). Fish should work with any POSIX-compliant UNIX providing it has a Bourne-compatible shell as /bin/sh and basic file manipulation utilities like cat, chgrp, chmod, chown, cp, dd, env, expr, grep, ls, mkdir, mv, rm, rmdir, sed and wc. If Perl is available, it is used instead. In that case, only env and /bin/sh are needed out of the list above, and using Perl has the advantage of being much faster. I've used fish to log in to Linux, FreeBSD, Mac OS X and Solaris with equal success.

Most of what is discussed in this article holds true of the ftp and sftp kioslaves as well, which present much the same interface to the user.

Here, I'd like to use fish to log in to my Apple Macintosh computer to grab some files and copy them locally. Because it's running a variant of UNIX with env and Perl and has sshd running, fish will work with it just fine. To log in to a remote host, type:

```
fish://username@host
```

into the Konqueror address bar. If your remote sshd is listening on a custom port, you can specify this at the end of the URL:

```
fish://username@host:9999
```

If this is the first time you've tried to connect to this host over SSH, Konqueror brings up a dialog noting that it cannot verify the authenticity of the host and asks for your input to proceed. Select Yes to indicate that you want to continue connecting to this host. SSH then caches the fingerprint of this host in your ~/.ssh/known\_hosts file.

Unless you have SSH keys set up for passwordless login, Konqueror shows a dialog asking you to enter your password. Ticking the box Keep password saves the password into your KWallet. This can be very convenient if you plan to access this host often.

Now I'm logged in and looking at my home directory on the remote host. I need to find the image I want out of the image files stored here. With the size of my digital photograph library, none of the images are named descriptively, and unfortunately, without being able to look at the images, I really have no idea which is the one I want.

One of the areas where fish really shines is the way it lets you preview remote files in much the same way you would preview local ones. KDE's preview mechanism is flexible and powerful, enabling you to see previews



Figure 1. KDE's ability to display previews over remote protocols makes finding the file you want a breeze.



www.asacomputers.com

1-800-REAL-PCS

## Hardware Systems For The Open Source Community—Since 1989

(Linux, FreeBSD, NetBSD, OpenBSD, Solaris, MS, etc.)

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

### AMD Opteron™ Value Server—\$795

- 1U 14.3" Deep
- AMD Opteron™ 140
- 512MB RAM *Max 8GB*
- 40GB IDE HDD
- 2x 10/100/1000 NIC
- Options: CD, FD or 2nd HD, RAID



### Front I/O Dual AMD Opteron™ Cluster Node—\$1,275

- 1U Dual AMD Opteron™ Capable Front I/O
- Single 240 AMD Opteron™
- 1GB RAM *Max RAM 16GB*
- 80GB HDD
- Dual PCI Expansion Slot



### 8 Hot Swap Bays in 2U AMD Opteron™—\$1,780

- 1 of 2 AMD Opteron™ 240
- 512MB RAM *Max 16GB*
- 3x80GB IDE RAID # 5
- 2xGigE, CD+FD
- Options: SATA/SCSI, Redundant PS



### No Frills AMD Opteron™ Storage Server—\$7,640

- 6TB+ IDE/SATA Storage in 5U
- Dual AMD Opteron™ 240
- 512MB RAM
- 6TB IDE Storage
- Dual GigE, CD
- Options: SATA HDD, DVD+RW etc.



### Your Custom Appliance Solution

Let us know your needs, we will get you a solution



### Custom Server, Storage, Cluster, etc. Solutions

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



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.




Figure 2. The Kate Open File Dialog. Displaying a Remote Location over Fish

of images, PDF files, fonts and office documents as well as listen to previews of audio files.

To enable previews over fish, select Configure Konqueror... from the Konqueror Settings menu. Navigate to Previews & Meta-Data in the dialog that appears. Under Internet Protocols, tick the box next to fish. Adjust the slider beneath Maximum file size to a sane value for the speed of the connection you're using. If you're accessing hosts over the Internet, you probably don't want Konqueror to attempt to generate thumbnail images of 100MB TIFF files. Click OK and you're done.

To fine-tune exactly the specific file types for which you want to see previews, look under Preview in the View menu. Here you can enable various file types for showing previews, and also quickly toggle between enabling and disabling previews globally. If you had a Konqueror window open while adjusting these settings, you may need to click reload before you see any changes.

Now I can see previews of my images, and I've found the one I want. I can drag and drop it to an open Krita window for editing and save my changes afterward directly to the remote host. Or, I simply could copy the file locally for further processing, also using drag and drop.

Fish lends itself well to the simple maintenance of remote Web sites. Tasks like moving, renaming and changing permissions on remote files are easy when you approach them from Konqueror's familiar file management interface.

Changing permissions for files on a Web server is a common task. Often the default permissions on files created on your Linux box can be too restrictive and result in a Forbidden error message when people try to access the files over the Web. To change the permissions of a file over fish, right-click on the file and select Properties. On the Permissions tab choose Can read for both Group and Other permissions, and click OK. If you were changing permissions for a CGI, you also would tick the box labeled is executable.

In KDE 3.5, Kate gained the ability to save a working set of documents as a session to be resumed easily again later. When coupled with fish, this makes, for me at least, the ultimate Web development environment.

To open up all of the HTML, CSS and other files associated with my Web site in Kate, I can drag and drop them from an open Konqueror window showing me a fish session to the remote host. Dragging and dropping multiple files in this way causes them to be opened as separate files that I can navigate between in the Documents pane.

If you prefer to use a more GUI-oriented Web development tool rather than editing files manually in Kate, the full-featured integrated Web development environment, Quanta Plus, is also KIO-enabled.

You can use the File→Open dialog to take advantage of KIO network transparency. All standard KDE file dialogs have support for kioslaves. To load files over fish, select Open from the File menu. In the location bar at

the top of the file dialog, type your fish URL:

`fish://canllaith@canllaith.org`

Navigate to where the files associated with your Web site or where other projects are kept, and open the files you want to be a part of the session. To open multiple files at once, hold down the Ctrl key and click once on each file you want to open. When you're finished selecting files, click Open.

If you plan to access this remote host often, you can add a shortcut to the icon bar on the left of the dialog. Choose the directory to which you want the shortcut to point, and drag it to the icon bar. You can edit the name, icon and other attributes of this shortcut by right-clicking on it and selecting Edit Entry.

Once Kate has loaded the files you want to save as a session, select Save As from the Sessions menu. Enter a name for your new session—as I'm saving various files relating to my Web site, I call my session canllaith.org.

In the future, when you want to load this set of remote files quickly, you can choose the saved session from the Sessions menu. I use this tool time and time again. At the end of a day of working on a Web site and opening various remote files from all over the directory tree, I simply save them as a session with the date and a short descriptive name. Next time, remembering where I left off and what files I was working on is a little bit easier.

## Samba

The smb kioslave included with KDE lets you browse Microsoft Windows smb file shares. It requires that you install libsmclient. If you navigate to smb:/ in Konqueror (or use the nifty Alt-F2 shortcut described below), you will be shown any Windows workgroups found, and you can browse through them for the host you want. You also can specify a host or a specific share of a host directly with:

`smb://username@hostname/share`

Like fish, if you don't specify a user name, Konqueror prompts you for a user name and password pair that you can save with KWallet. If you always use a particular user name/password pair on your Samba network, rather than having to save passwords individually for every host you access, you can configure this to be supplied automatically by KDE. In the KDE Control Center, navigate to Internet & Network→Local Network Browsing. Here you can enter the default user name and password pair you want KDE to use for its Samba client.

As well as adding shortcuts to the File dialog, you also can add desktop shortcuts to hosts you want to access frequently. To create a desktop shortcut to an smb URL, right-click on the KDE desktop and select Create New→Link to Location (URL)... from the context menu. Fill in the smb:// (or fish) URL to the share to which you want to create the shortcut in the box labeled Enter link to location (URL):. KDE fills in the filename box with a suitable name, or you can choose your own. Click OK and you're done.

As well as accessing kioslaves through the Konqueror address bar and KDE standard file dialogs, you can load kioslaves quickly with the KDE Run



Figure 3. Creating a Desktop Shortcut for a Location Accessible over Samba




Figure 4. The locate kioslave shows the results of the query “kio” in my home directory.

Command box. Try pressing Alt-F2 to bring up the Run Command box, and type `help:/kwrite`. A Konqueror window is launched showing you the KWrite Handbook. This works with all kioslaves and is a handy way of looking up help pages or loading a remote URL quickly, if, like me, you tend to have a rather cluttered screen.

Many other interesting kioslaves are included with KDE, and you can download other third-party efforts from [kde-apps.org](http://kde-apps.org) as source code that can be compiled against a recent KDE version. To find them, search for “kio” on the KDE-apps.org search page. If you want to compile the kioslaves you’ve downloaded, you need to have a working C++ compiler and the appropriate development libraries for KDE and Qt installed. Usually these are packaged separately from the KDE runtime libraries.

To find out which kioslaves you have installed, type `help:/kioslave` in the Run Command box or the Konqueror address bar. This is the KDE help kioslave, which lets you access the help documentation for installed KDE programs through Konqueror. Some of the more interesting kioslaves include:

- `cgi`: this kioslave executes CGI programs without needing to have a running Web server. It is really handy for off-line local testing of CGI scripts.
- `locate`: Ubuntu includes `kio-locate` by default, and you can download the sources for other distributions from [kde-apps.org](http://kde-apps.org). `kio-locate` is a kioslave for `locate` or `slocate`. Typing `locate: query term` into any KIO-enabled field displays the results from the `locate` database. This is immensely convenient when combined with the File dialog. Want to open that budget spreadsheet in KSpread, but you realise you can’t quite remember where you saved it until after you’ve launched the application? Without having to leave the File dialog, `locate:/` comes to the rescue.

- `tar`: this kioslave allows you to browse the contents of tar, tar.bz2 and tar.gz archives. It’s registered as the default handler for these files within KDE. This lets every KDE application handle loading and saving files to archives transparently without needing to extract them. With previews enabled, it’s easy to find the single file that you want out of the hundreds or even thousands in the archive.
- `zip`: this kioslave lets you browse the contents of zip archives, much like the tar kioslave does for tar archives.
- `info/man`: the info and man kioslaves provide a friendly interface to reading man and info pages. The info kioslave in particular makes navigating pages much easier with a mouse-driven browser interface that’s more simple to use than the command-line tool.
- `audiocd`: this kioslave provides a simple interface for ripping and encoding files from music CDs to Ogg, MP3 or flac using drag and drop.

Konqueror is an application with amazing flexibility as both a Web browser and file manager, due mostly to its extensibility with kioslaves. The kioslaves featured above are barely the tip of the iceberg. Experiment with those listed in `help:/kioslave` to see what else Konqueror can do. ■

Jes Hall is a KDE developer from New Zealand who is passionate about helping open-source software bring life-changing information and tools to those who would otherwise not have them. She welcomes comments sent to [jhall@kde.org](mailto:jhall@kde.org).

# Linux Laptops

Starting at \$990

---

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





**309.34.CUBED**

[shoprcubed.com](http://shoprcubed.com)

# YELLOW DOG LINUX

## INSTALLS NEATLY ON AN iPod

Forget bootable USB pendrives and use an iPod to boot Linux on a Mac.

Dave Taylor



**THE CONCEPT'S GREAT:** what would it be like to have a pocket-size device that I could plug in to just about any Macintosh and by simply rebooting the computer be running a full-blown Linux installation? There are oodles of Linux OSes for Intel architectures, of course, but the Mac, until very recently, has been built around the Motorola architecture, so the number of choices are rather fewer.

One of the few Linux OSes for the PowerPC is called Yellow Dog, from Terra Soft Solutions, [www.yellowdoglinux.com](http://www.yellowdoglinux.com). It costs about \$60 US for the install CDs and documentation or \$30 US for the "geek edition" (that's just the install CDs), or you can download it for free from the Web site. And, let me answer the obvious question: because Mac OS X already is a UNIX (basically FreeBSD with lots of added stuff, much of which you can find in Darwin, [www.apple.com/darwin](http://www.apple.com/darwin)), why bother with a Mac Linux? The answer is that although Mac OS X is a splendid mating of a UNIX operating system with all the graphical goodness of Apple's user interface design, it's still not Linux. If you're in a Linux environment and want to run KDE or GNOME, you don't have to graft it onto Mac OS X if you can run a Linux designed for the Mac platform instead. Besides, isn't it kinda cool anyway?

Anyway, I had a spare Apple iPod, a first-generation 5GB device that worked via the FireWire interface rather than the more modern USB connection, and I was assured by the folks at Yellow Dog that I could squeeze YDL into as small as 1GB. I have plenty of space on a 5GB device. Of course, I already had a gig of music and audio books I wanted to preserve, so the first test was to see if I could repartition the device to grab 3GB for Linux and keep 2GB for audio and iPod content. The perfect stealth Linux device, right?

So, one afternoon I decided to take the plunge and hooked up my iPod to my PowerBook computer and inserted the first of the YDL 4.1 install disks and restarted the Mac, holding down the C key to force the device to boot off the CD-ROM, not the internal hard disk. When prompted, I typed in `install firewire` and away we went.

### Partitioning the iPod for YDL

New to the 4.x version of Yellow Dog is the inclusion of the popular Anaconda graphical installer, which makes everything quite a bit easier. It lets you resize existing drive partitions to make space for the new operating system. The new partitions also can be made bootable, which is a critical component for the success of this project.

Theoretically, partitioning should be pretty easy. I have a 5GB iPod FireWire device and am using just a wee bit more than 1GB of it for music. I'll resize the iPod drive to 2GB and have 3GB spare to repartition as an ext3 filesystem and be good to go.

Well, that's the theory, but it doesn't quite work out that way.

Part way through the install process, managed by Anaconda, I have the option of accepting an automatic partitioning scheme or using Disk Druid to work with my disk partitions manually. I take the latter path and am glad to see that one of the drives is identified as "Drive /dev/sda (4769MB) Model: Apple iPod", so there's no worry that I'll accidentally reformat or resize my laptop

drive, which would be quite ungood. To resize the iPod drive, I simply choose that partition and click Edit in the Disk Druid, and then specify that I want it to be 2,000MB rather than 4,769MB (which should give me 2.7GB for Linux). It promptly recalculates that to be 1,999MB and within about 90 seconds rebuilds the iPod disk partition, leaving a big chunk of space unallocated.

Here's where I get into trouble, because I'm a UNIX geek who is sure that I can proceed without reading any darn manual or instructions. Yeah, even Terra Soft expects this and has a note in the installation guide (which I didn't read until afterward, of course) saying, "User error is common. Not because people lack intelligence, but because people are smart and too determined to jump into their new operating system without reading the Guide to Installation. Especially those of you who are Linux Experts—you know who you are!" Yeah, yeah, yeah.

Fortunately, the trouble ate up only time, and didn't corrupt anything. Basically, although I figured that I simply could create one partition that was all the available space, Disk Druid wouldn't let me proceed without also creating an Apple Boot partition, and then, after I figured that out (the Apple Boot partition is instead of ext3, and not the same as the /boot mountpoint for an ext3 partition), it also insisted I create a swap partition too.

More than once it complained, and I had to back up and resize the new partition down, then create an additional partition, but, finally, here's where I ended up (Table 1).

Table 1. Partition Breakdown

Partition	Size	File Type	Mountpoint
sda3	2000MB	foreign	
sda4	1MB	Apple Bootable	
sda5	125MB	swap	
sda6	2596MB	ext3	/

If you're paying attention, you'll see that the swap space is really too small. You should have at least the same swap space as your physical memory, and typically 1.5x is a better size for performance reasons. Because I have 756MB of RAM, that means I should have at least a 756MB swap space. Oh well. I indicated that I was okay with a nonrecommended size and proceeded anyway.

Elapsed time: 1 hour.

## Basic Network Configuration and Installation

After dancing the Disk Druid dance for almost an hour, it was a distinct pleasure to get to a prompt asking about DHCP and network configuration. I picked all the basic defaults, except I skipped configuring a firewall. It didn't like that, but let me proceed after giving me a little lecture on system security.

As I originally chose a Personal Workstation configuration, it meant that my default package set was X Window System + KDE + OpenOffice.org + Mozilla + Evolution + IM tools + games. Not good. Why? Because my disk partition was 367MB too small.

Going back to the proverbial drawing board, I started trying to pull out individual applications, guessing how much each one would take of the installation. It's amazing, really, after all these years, that Anaconda doesn't indicate how big each package is when you're trying to navigate through it. Instead, I piddled around removing Gaim (a multi-IM utility) to save 41MB; XChat (an IRC client) to save 5MB; all the sound and video applications (saving 57MB); all the graphics applications, including The GIMP and ImageMagick (saving 100MB); and the KDE component kdeggraphics (saving 26MB). I attempted to re-install, and wouldn't you know it—still too big, by 185MB.

As you might expect, this was pretty tedious. But when I dug around in the Office Utilities area, I was amazed and delighted to see that the support package openoffice.org-18n (a package with lots of localization libraries for OpenOffice.org) was a whopping 668MB in

size. Because I didn't envision that I'd be editing documents in German, Hebrew or Kanji, I happily deleted it and re-added all the individual apps I'd deleted earlier. I even threw in kdegames, eating up 23MB, but hey, who doesn't like games?

Finally, 75 minutes after I started the process, I actually was able to proceed with the full installation. It took 18 minutes before I saw "installation finished", which I attribute to the fact that the iPod FireWire drive is slower to access than the internal hard drive in the PowerBook.

## Reboot and Be Happy

I held down the OPTION key on the keyboard during the boot sequence to be able to access the Yellow Dog Linux OS as an alternative to the Mac OS X on my main PowerBook drive. After about 60 seconds of hunting for options, it showed me both Mac OS X Tiger and Yellow Dog Linux. Eureka!

I selected YDL, clicked on the continue button (an arrow) and then was in the yaboot program, where I pressed L for Linux and sat back. Lots of status information scrolled past, including the information that eth0 (the built-in Ethernet port) failed to initialize, which made sense as I wasn't hooked up to a network. Otherwise, I was soon looking at the attractive KDE login window, to which I typed in my new user account information that I'd specified seconds earlier in the first boot utility.

I then was prompted to select display specifics and was pleased to see that one of the display manufacturers listed was Apple. Scrolling down the long, detailed list, I found the right match: "Apple Titanium PowerBook G4" and accepted the defaults for that display.

The next step was particularly satisfying, as it asked about audio hardware configuration and worked with the default settings. Previously, when I had installed an earlier version of YDL on the PowerBook, the audio subsystem had failed, never to work again—a valuable upgrade by itself.

Once the setup was done for KDE, I was running in a full-blown Linux/KDE environment, with all the applications, utilities and games I could want. It was fast, smooth and quite a delight to have a different desktop and user environment on my system.

But, I wanted to test and ensure that everything still worked properly, so I shut down YDL, and sat looking at a dark screen, realizing that there was really no way to know when it had completed its shutdown. Fortunately, I also was watching the iPod screen, and once the system finished shutting down, the iPod switched from "do not disconnect" to an Apple logo, and then rebooted into iPod mode.

Indeed, the iPod works perfectly. All my audio files remained intact, and now when I go to the System Information area on the iPod, it shows that the storage capacity of the unit is 1.96GB rather than the earlier 5GB value. Perfect!

Everything unplugged, I restarted the PowerBook and was gratified to watch it quickly and easily restart in Mac OS X, without any indication that I'd installed anything unusual, touched any hard drives or restarted in a foreign OS just a few minutes earlier.

Success!

## And in the End

Alright, it's geeky, but I think it's way cool to have an iPod that can boot any G4 Mac into a full Linux work environment with only a few keystrokes. If you need Linux functionality and don't want to touch your existing Mac OS X systems, this can be a great solution, and you don't even lose the functionality of your iPod along the way. Indeed, a quick search on eBay shows that you can pick up one of these ancient 5GB iPod units for less than \$60 US, on average.

There are some caveats about this installation, however, particularly regarding the very latest iPod systems, which have a slightly different filesystem. If you are going to proceed with this, don't follow my lead but start on the Terra Soft site and read the hardware and configuration notes. It'll save you a lot of heartache down the road. ■

---

Dave Taylor has been involved with the UNIX community since 1980 and was the original author of The Elm Mail System. He's written 20 books, including *Teach Yourself Unix in 24 Hours* and *Wicked Cool Shell Scripts*. He invites all true Linux fans to visit his Weblog at [www.askdaveataylor.com](http://www.askdaveataylor.com).



ILLUSTRATION BY JEFFREY WONG

## Who needs NFS or Samba when you can mount filesystems with SSHFS?

Matthew E. Hoskins

**Tools like** `scp`, `sftp` and `rsync` allow us to copy files easily and securely between these accounts. But, what if we don't want to copy the files to our local system before using them? Normally, this would be a good place for a traditional network filesystem, such as NFS, OpenAFS or Samba. Unfortunately, setting up these network filesystems requires administrator access on both systems.

### **Enter SSHFS and FUSE**

Luckily, as long as you have SSH access, you can use SSHFS to mount and use remote directory trees as if they were local. SSHFS requires no special software on the remote side, just a modern SSH server with support for the SFTP extension. All modern Linux distributions support this extension, which was added to OpenSSH in version 2.3.0, released in November 2000.

SSHFS is built upon the FUSE user-space filesystem framework project. FUSE allows user-space software, SSH in this case, to present a virtual filesystem interface to the user. This is similar to how the `/proc` and `/sys` filesystems present kernel internals in the form of files in a directory tree. SSHFS connects to the remote system and does all the necessary operations to provide the look and feel of a regular filesystem interface for remote files.

## Installing SSHFS and FUSE

I am using Fedora Core 4 on the workstation where I will be mounting the remote directories. The remote system is also Fedora Core 4, but any \*NIX system running a modern SSH server with the SFTP extension will work. Your Linux kernel also must have the user-space filesystems feature enabled, either built-in or as a module.

All the software required for SSHFS is available as packages installable with yum for Fedora Core 4. Simply run:

```
$ yum install fuse-sshfs
```

This installs SSHFS, FUSE and the fuse-lib dependencies automatically. You also can build SSHFS from source, but more on that later.

Before you can use SSHFS or any other FUSE-based filesystem as a nonroot user, you must first add those users to the fuse group. In my case, my user name is matt. This can be done from the command line as root with the command:

```
$ usermod -a -G fuse matt
```

The fuse group lets you limit which users are allowed to use FUSE-based filesystems. This is important because FUSE installs setuid programs, which always carry security implications. On a highly secured system, access to such programs should be evaluated and controlled.

After installing and configuring the software, we are ready to give it a whirl. To demonstrate the basic functionality, I will make a connection to a remote system called my.randombox.com. The default operation for SSHFS is to mount the remote user's home directory; this is the most common use of SSHFS. Just like mounting any other filesystem, you need an empty directory called a mountpoint under which the remote filesystem will be mounted. I create a mountpoint named randombox\_home, and then invoke the sshfs command to mount the remote filesystem. Here is how it's done:

```
$ cd $HOME
$ mkdir randombox_home
$ sshfs matt@my.randombox.com: randombox_home
matt@my.randombox.com's password: *****
$ ls -l randombox_home/
-rw-r----- 1 matt users 7286 Feb 11 08:59 sshfs.article.main.txt
drwx----- 1 matt users 2048 Mar 21 2001 projects
drwx----- 1 matt users 2048 Dec 1 2000 Mail
drwxr-xr-x 1 matt users 4096 Jun 8 2002 public_html
$ cp ~/my_web_site/index.html randombox_home/public_html/
```

That's it. My home directory on my.randombox.com is now mounted under the directory randombox\_home on my local workstation. In the background, FUSE, SSHFS and the remote SSH server are doing all the legwork to make my remote home directory appear just like any other local filesystem. If you want to mount a directory other than your

home directory, simply put it after the colon on the sshfs command line. You even can specify / to mount an entire remote system. You will, of course, have access only to the files and directories for which the remote user account has permission. Everything else will get "Permission denied" messages. An example of this is shown below:

```
$ cd $HOME
$ mkdir randombox_slash
$ sshfs matt@my.randombox.com:/ randombox_slash
matt@my.randombox.com's password:
$ ls -l randombox_slash/
total 0
drwxr-xr-x 1 root root 4096 Nov 15 10:51 bin
drwxr-xr-x 1 root root 1024 Nov 16 07:11 boot
drwxr-xr-x 1 root root 118784 Jan 26 08:35 dev
drwxr-xr-x 1 root root 4096 Feb 17 10:37 etc
drwxr-xr-x 1 root root 4096 Nov 29 09:30 home
drwxr-xr-x 1 root root 4096 Jan 24 2003 initrd
drwxr-xr-x 1 root root 4096 Nov 15 10:53 lib
drwx----- 1 root root 16384 Nov 11 10:21 lost+found
drwxr-xr-x 1 root root 4096 Nov 11 10:38 mnt
drwxr-xr-x 1 root root 4096 Jan 24 2003 opt
dr-xr-xr-x 1 root root 0 Jan 26 08:11 proc
drwx----- 1 root root 4096 Mar 3 09:34 root
drwxr-xr-x 1 root root 8192 Nov 15 13:50 sbin
drwxrwxrwt 1 root root 4096 Mar 5 18:41 tmp
drwxr-xr-x 1 root root 4096 Nov 11 10:55 usr
drwxr-xr-x 1 root root 4096 Jan 20 08:16 var
$ cat randombox_slash/etc/shadow
cat: randombox_slash/etc/shadow: Permission denied
$ ls -l randombox_slash/root/
ls: reading directory randombox_slash/root/: Permission denied
total 0
$ ls -l randombox_slash/home/matt/
-rw-r----- 1 matt users 7286 Feb 11 08:59 sshfs.article.main.txt
drwx----- 1 matt users 2048 Mar 21 2001 projects
drwx----- 1 matt users 2048 Dec 1 2000 Mail
drwxr-xr-x 1 matt users 4096 Jun 8 2002 public_html
$
```

## Automating the Connection

As you can see from the above examples, I needed to type my password to complete the SSH connection to the remote system. This can be eliminated by creating a trust relationship between the local and remote user accounts. This is not appropriate in all situations, because it essentially makes the accounts equivalent from a security perspective. Any malicious activity on one account can spread to other systems via the trust, so take caution and fully understand the implications of setting up trust relationships. To begin setting this up, you need to create an SSH key pair, which consists of public and private key files named `id_rsa` and `id_rsa.pub`, respectively.

The public key is copied to the remote system and placed in the `$HOME/.ssh/authorized_keys` file. Some systems may use the filename `authorized_keys2` in addition to or instead of `authorized_keys`.

This allows any user in possession of the private key to authenticate without a password. We create the key pair using the command `ssh-keygen`. The files are placed in the proper locations automatically on the local system in the




Figure 1. Set up a GNOME startup command to mount an SSHFS share.




Figure 2. GNOME automatically mounts the remote directory.

`$HOME/.ssh` directory. Because we already have my remote home directory mounted, appending the public key to the `authorized_keys` file is extra easy. Below are all the steps required (assuming you created the equivalent of the `randombox_home` directory and mounted it):

```
$ cd $HOME
$ ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/home/matt/.ssh/id_rsa): <ENTER>
```

```
Enter passphrase (empty for no passphrase): <ENTER>
Enter same passphrase again: <ENTER>
Your identification has been saved in /home/matt/.ssh/id_rsa.
Your public key has been saved in /home/matt/.ssh/id_rsa.pub.
The key fingerprint is:
fa:e7:7c:e1:cb:7b:66:8b:67:07:05:99:7f:05:b9:4a matt@myworkstation
$ mkdir randombox_home/.ssh
$ chmod 700 randombox_home/.ssh
$ cat .ssh/id_rsa.pub >> randombox_home/.ssh/authorized_keys
$ chmod 600 randombox_home/.ssh/authorized_keys
```

In the above example, we create the key pair with an empty passphrase, then append the public key to the `authorized_keys` file in the remote home directory and set the permissions. After this is done, I no longer need to type the password when connecting to the remote account. To test this, first we unmount the remote home directory with the following command:

```
$ fusermount -u randombox_home
```

To test the trust relationship, we can try to run the `uptime` command on the remote system:

```
$ ssh matt@my.randombox.com uptime
12:20:40 up 38 days, 4:12, 0 users, load average: 0.11, 0.04, 0.01
```

Good, no password needed. The trust relationship is working properly. If you have trouble getting this trust relationship to work, check the permissions on the files in `.ssh` on both systems. Many times lax permissions prevent SSH from using key files. Also, take a look at the `syslog` log files. OpenSSH's `sshd` server logs messages into `syslog`, which often are helpful in diagnosing key file problems. You may have to increase the logging verbosity level in the `sshd_config` file, usually found in `/etc/ssh/`.

You also can debug the connection by running `ssh` in the above example with the `-vvv` option to turn up verbosity. Now, let's mount the remote directory again. This time it does not prompt for my password:

```
$ cd $HOME
$ mkdir randombox_home
$ sshfs matt@my.randombox.com: randombox_home
$ ls -l randombox_home/
-rw-r----- 1 matt users 7286 Feb 11 10:33 sshfs.article.main.txt
drwx----- 1 matt users 2048 Mar 21 2001 projects
drwx----- 1 matt users 2048 Dec 1 2000 Mail
drwxr-xr-x 1 matt users 4096 Jun 8 2002 public_html
```

## Integrating with the GNOME Desktop

In the last example, we configured and automated non-interactive mounting of a remote directory. Because we're no longer being prompted for a password, we can integrate SSHFS mounting into scripts, or better yet the GNOME desktop. To configure GNOME to mount our remote home directory automatically, we configure the SSHFS mount command as a session startup program. This is done from inside the Sessions preferences dialog. Navigate to `Desktop`→`Preferences`→`More Preferences`→`Sessions`→`Add`, and fill in the dialog as shown in Figure 1.

Upon the next login, GNOME automatically mounts the remote directory for me, as you can see in Figure 2.

Note that GNOME does not reliably kill this command upon exiting the session. You can unmount the remote directory manually using the `fusermount -u randombox_home` command. Another option is to automate the unmount by modifying the `$HOME/.Xclients-default` file to run the `fusermount` command as follows:

```
#!/bin/bash
# (c) 2001 Red Hat, Inc.

WM="gnome-session"
WMPATH="/usr/bin /usr/X11R6/bin /usr/local/bin"

# Kludged to run fusermount upon gnome logout. 20060301-MEH
for p in $WMPATH ; do
    [ -x $p/$WM ] && $p/$WM; fusermount -u randombox_home; exit 0
done

exit 1
```

Be aware that the `.Xclients-default` file is rewritten every time you run the `switchdesk` utility. You have to modify this file every time you use the `switchdesk` utility to change your default desktop windowing manager.

Finally, you can add the appropriate `sshfs` commands in the boot startup file that is appropriate for your distribution. This way, your system will mount all the `SSHFS` directories automatically each time you boot your desktop.

### Building SSHFS from Source

If your particular Linux distribution does not prepackage `SSHFS`, or if you simply want to build it from source, this also is pretty easy. First, confirm that you have installed whatever files or packages are required for kernel module development. You need these to build the `FUSE` kernel module. Then, download the latest source tarballs for both `FUSE` and `SSHFS` from SourceForge (see the on-line Resources). Place the downloaded tarball files in a temporary directory, then build and install using the following commands in that directory:

```
$ tar -xzf fuse-2.5.2.tar.gz
$ cd fuse-2.5.2
$ ./configure --prefix=/usr
$ make
$ su -c "make install"
$ cd ..
$ tar -xzf sshfs-fuse-1.5.tar.gz
$ cd sshfs-fuse-1.5
$ ./configure --prefix=/usr
$ make
$ su -c "make install"
```

After everything is installed, you are ready to perform any of the examples presented previously. After installation, the `sshfs` and `fusermount` commands are installed in `/usr/bin`.

### Conclusion

`SSHFS` and `FUSE` allow any remote storage to be mounted and used just like any other filesystem. If you can log in with `SSH`, you have all the access you need.

As I said earlier, `FUSE` is a framework for creating user-space filesystems. `SSHFS` is only the tip of the iceberg. There are `FUSE`-based filesystems to encrypt your files (`EncFS`) transparently, browse Bluetooth devices (`BTFS`) or mount a CVS repository as a filesystem (`CvsFS`). Perhaps you were wondering what to do with all that free space in your Gmail account? Well, `GmailFS` allows you to mount your Gmail account and use it like a local filesystem. See the `FUSE` Web site for these and more great projects, or perhaps you would like to write your own. `FUSE` has language bindings for Perl, Python, TCL, C, C#, Ruby and others. ■

Resources for this article: [www.linuxjournal.com/article/8943](http://www.linuxjournal.com/article/8943).

Matthew E. Hoskins is a Senior UNIX System Administrator for The New Jersey Institute of Technology where he maintains many of the corporate administrative systems. He enjoys trying to get wildly different systems and software working together, usually with a thin layer of Perl (locally known as "MattGlue"). When not hacking systems, he often can be found hacking in the kitchen. Matt is a member of the Society of Professional Journalists. He is eager to hear your feedback and can be reached at [matt@njit.edu](mailto:matt@njit.edu).

## carinet cares about your Dedicated Server budget!

Our unbeatable pricing helped save over 3000 piggy-banks last year alone!

Better Servers. Better Service.

**\$59** /mo  
Celeron D 3GHz  
1GB RAM, 160GB SATA, 1200 GB/month

**\$139** /mo  
Xeon HT 3GHz  
2GB RAM, 160GB SATA  
1200 GB/mo

Upgradable to Dual HT!

**carinet**  
Cari.net/LJ or 888.221.5902

## An Introduction to Gambas

If you think Visual Basic is almost useful, here's a way to almost get it on Linux.

MARK ALEXANDER BAIN

**Have you ever** wanted to use Visual Basic on Linux? Why? Well, you could be like me and have spent many years programming in VB, but want to move to Linux without having to learn a new language. It could be that you're brand new to programming and need something you can learn quickly and easily—and still be able to produce a good quality application.

Well, now you can—almost.

Gambas is short for Gambas is—almost—Basic, and it has been designed to look at the good things VB can do for Windows and then does them for Linux. Above all, Gambas is easy to use—as this article shows. I explain how to build a useful application in Gambas—a bug-tracking system that stores its information in a MySQL database.

Installation is simple. First go to the Gambas Web site, and check the Distributions & OS page—this is just to make sure there are no known peculiarities with your flavour of Linux. Then, go to the Download page and get the most current, stable version (1.0.9 at the time of this writing). If you've done this type of thing before, simply carry on and get yourself ready to use Gambas; if not, don't worry—we're nearly there.

Open a terminal and move to the directory where you've saved the bz2 file. If you're going to use 1.0.9, it will be called `gambas-1.0.9.tar.bz2`. Now `bunzip2` the file, and follow the installation instructions (unless the distribution page has given you some additional parameters for your distribution).

With that, you're ready to use Gambas. Type `gambas` on the command line, and the welcome screen appears (Figure 1).

The Gambas screen gives you access to example applications—you'll find these very useful if you are new to programming. In this case, click on New Project. Gambas now displays its project creation wizard, so follow the instructions to create your first project. When asked, create a graphical project, set the name to `bugTracker` (note that underscores are not allowed), and then set the title to Bug Tracker. You also will be asked where to store

your project. I suggest you create a new directory called Gambas in your home directory and then use it for all future projects as well.

Next, we jump straight in to the creation of a new form. Right-click in the project window, select New and then Form. Gambas puts you into its form creation wizard. All you need to do now is give the form a name—call it `frmBugTracker`. (Don't leave it as `Form1`. That's *very bad practice*.)

Now, you can start adding the elements to the form—and, the first one to add is a Close button. Why do this first? Quite simply, you always want to be able to close a form, or an application, for that matter, cleanly and easily, so get into this habit as quickly as possible. To create the button, click on its icon in the toolbox (the icon is a box with OK on it), and then use the left-mouse button to draw it onto the form. By default, the button is called `Button1`, but exactly like the form, we rename it. Click on the button, and press F4 to display the Properties box. Change its name to `btnClose` and its text to `Close`.

The button won't do anything yet—we have to add some code to it, which is really, really easy. If you double-click on the button, Gambas takes you to the code window, and you'll find an empty `btnClose_Click` subroutine. Modify it so that it says:

```
PUBLIC SUB btnClose_Click()

    ME.Close

END
```

You should notice something as you type in the code—as soon as



Figure 1. The Gambas Welcome Screen




Figure 2. Designing a New Gambas Form



~~I CAN'T GET MY DATA OUT OF OUR CRM SYSTEM.  
CUSTOMIZING OUR CRM SYSTEM IS OFF-LIMITS.  
OUR CRM SYSTEM IS SO INFLEXIBLE.~~

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

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

Get hooked on SugarCRM. Visit: [www.sugarcrm.com/linuxjournal](http://www.sugarcrm.com/linuxjournal) or call +1 408.454.6941.

Commercial Open Source Customer Relationship Management



**SUGARCRM**  
COMMERCIAL OPEN SOURCE  
[www.sugarcrm.com](http://www.sugarcrm.com) • [www.sugarforge.org](http://www.sugarforge.org)  
1.875.SUGARCRM • +1.408.454.6941

you put in the period, a drop-down menu appears, giving you all of the methods and parameters associated with the ME object—in this case, ME refers to the form, so ME.Close means close this form.

Suppose you want to see the results of your hard work now. Go to the Project window and click on the little green run button. And there you are—your first Gambas application. You even can close the form with your brand-new button.

Before building the application itself, we need to think about what we want the bug tracker to do. Then, we need to have a look at the data—how it is to be arranged, and what is going to be stored.

The bug tracker will have to do the following:

- Have the details of new bugs entered.
- Record who raised the bug.
- Have a bug assigned to a programmer.
- Update the status of the bug.
- Record when the bug was raised.
- Record the release for the bug fix.
- Provide the ability to view new, working and complete bugs.

The data required is therefore:

- Who raised the bug.
- Who is fixing the bug.
- Bug details.
- Developer details: ID, first name, surname and user name.
- Bug Details: ID, description, date created, ID of raiser, ID of developer, status, release number and application details.

From this, we can start building a database schema. Start by creating a file (such as database\_schema.sql) that we will use to create the database:

```
/*
First you must create the database. The listing
table includes only the user ids for "raised by"
and "coder".
*/
create database bugtracker;

create table bugtracker.listing (
id int auto_increment primary key,
details longtext,
application_id int,
release float,
raised_by_id int,
coder_id int,
status_id int,
created datetime,
priority int,
status int);

/*
The coder table is simple but includes the user name.
*/
```

```
create table bugtracker.coder (
id int auto_increment primary key,
surname varchar(50),
firstname varchar(50),
username varchar(50));
/*
Finally you can create reference tables for storing
application names and status titles.
*/
create table bugtracker.application (
id int auto_increment primary key,
name varchar(50));

create table bugtracker.status (
id int auto_increment primary key,
title varchar(50));

/*
With the tables created you can add a user account to
the database for the bugTracker application to log
on to.
*/
GRANT select,insert,delete,update ON bugtracker.*
TO bugtracker@localhost IDENTIFIED BY 'mypassword';

/*
Next you can start loading the data that will make
the application work. The key information is the coder data.
*/

insert into bugtracker.coder (username,surname,firstname)
values ('bain','bain','mark');
/*
Finally add some dummy data so that you can see
the application working as soon as possible
*/
insert into bugtracker.application (name)
values ('bugtracker');
insert into bugtracker.status (title)
values ('new');
insert into bugtracker.status (title)
values ('worked on');
insert into bugtracker.status (title)
values ('rejected');
insert into bugtracker.status (title)
values ('completed');
```

Create the database by typing:

```
mysql -uroot -p<root password> mysql < database_schema.sql
```


So, with the database in place, it's time to go back to Gambas to do some actual programming.

Now that we've built and loaded data into the database, the first thing to do is to connect to it, so we can start communicating with it. Gambas makes this very, very easy for us. It comes with components, and all we have to do is tell the application to use the appropriate component for connecting to databases. We will have to do a little coding as well, but that is easy too.

Go to the Project window, click on Project and then Properties. The Properties screen appears—click on the Components tab, and then check gb.db (Figure 3).

Gambas now can use the component to communicate with the

# X Marks the Slow Node!



## MPI Link-Checker™ to the Rescue!

A single slow node or intermittent link can cut the speed of MPI applications by half. Whether you use GigE, Myrinet, Quadrics, InfiniBand or InfiniPath HTX, there is only one choice for monitoring and debugging your cluster of SMP nodes: Microway's MPI Link-Checker™.

This unique diagnostic tool uses an end-to-end stress test to find problems with cables, processors, BIOS's, PCI buses, NIC's, switches, and even MPI itself! It provides instant details on how latency and bandwidth vary with packet size. It also provides ancillary data on inter-process and intra-CPU latency, and includes FastCheck!, which runs in CLI mode and checks up to 100 nodes per second.

**A complimentary one year license for MPI Link-Checker™ is installed on every Opreton based Microway cluster purchased in 2006.**

Wondering what's wrong with your cluster's performance, or need help designing your next one? Microway designs award-winning single and dual core AMD Opreton based clusters. Dual core enables users to increase computing capacity without increasing power requirements, thereby providing the best performance per watt. Configurations include 1U, 2U, and our 4U QuadPuter™ RuggedRack™—available with four or eight dual core Opretons, offering the perfect balance between performance and density.

Microway has been an innovator in HPC since 1982. We have thousands of happy customers in HPC, Energy, Enterprise and Life Science markets. Isn't it time you became one?

Call us first at 508-746-7341 for quotes and benchmarking services. Find technical information, testimonials, and newsletter at [www.microway.com](http://www.microway.com).



▲ Microway® Quad Opreton™ Cluster with 36 Opreton 880s, redundant power, 45 hard drives and Myrinet™ in our CoolRak™ cabinet.






Figure 3. Enabling the Data Access Component

database. All we have to do is add a little bit of code, and we do this in a module, which is a file for storing code so that it is usable by an application in general and not just a single form. For instance, the Close function we already created was only for the one form. We want to create code that is accessible for any forms that we create.

Right-click on the Project screen, click New and then Module. Gimp shows the Module Creation wizard. Just like the form and the button we made earlier, don't leave the module name as Module1. Call it something useful. In this case, call it Data. Once you create it, you can start adding code to it. Create the variables that will be used for database access:

```
PRIVATE myDB AS NEW Connection PUBLIC tmpRec AS Result
```

If you've used VB, you'll be happy with the format. If you're new to this, just take note that the PUBLIC statement makes the variable global—accessible to the whole application. If you don't want it to be available to the whole application, define it as PRIVATE.

The next (public) function makes the connection:

```
PUBLIC FUNCTION connect2db() AS Boolean
  WITH myDB
    .Type = "mysql"
    .Host = "localhost"
    .Login = "bugtracker"
    .Password = "mypassword"
    .Name = "bugtracker"
  END WITH
  TRY myDB.Open
  IF ERROR THEN
    Message ("Cannot Open bugtracker:" & Error.Text)
    RETURN FALSE
  END IF
  RETURN TRUE
END

PUBLIC FUNCTION Exec(sql AS String) AS Result
  RETURN myDB.Exec(sql)
END
```

We can call these functions from any form that we create. In this

case, we call them from frmBugTracker. Double-click on frmBugTracker in the Project window, and then double-click anywhere on the form itself. This takes you into the code window, and you should see:

```
PUBLIC SUB Form_Open()
END
```

Now we add code to tell the form to connect to the database, and we also add a function to carry out a simple security check:

```
PUBLIC SUB Form_Open()

  Data.connect2db

  IF (check_id() = FALSE) THEN
    message ("Unable to log on as " & system.user)
    ME.close
  END IF
END

PRIVATE FUNCTION check_id () AS Boolean
  data.Exec("select id " &
    " from coder" &
    " where username='" & system.user & "'")
  IF (data.tmpRec.Available ) THEN
    RETURN TRUE
  ELSE
    RETURN FALSE
  END IF
END
```

If you run the project now, little will have changed, apart from the fact that it will take a little longer to load—it now has to connect to the database. However, the form will check the user's Linux user ID against the list of coders on the database using the function check\_id. It displays a message and then closes the form if the ID is missing.

Next (keeping it simple), we create a pair of combo-boxes. One (cmdBugId) displays the list of bug IDs assigned to the current user. The other (cmdStatus) displays a list of the possible statuses. We then add a sub-route (loadCombos) to fill in the details of the combo-boxes. Once you have added the combo-boxes from the toolkit write the required subroutines:

Add this to the Data module:

```
PUBLIC SUB loadCombo (combo AS ComboBox,
  sql AS String)

  combo.Clear
  tmpRec = myDB.Exec(sql)
  FOR EACH tmpRec
    combo.Add (tmpRec[0])
  NEXT
END
```

Add this to frmBugTracker:

```
PRIVATE SUB loadCombos ()

  data.loadCombo(cmbBugid,"select l.id" &
    " from listing l, coder c" &
    " where l.coder_id=c.id" &
    " AND c.username='" & system.User & "'")

  data.loadCombo(cmbStatus,"select title from status")
END
```




Figure 4. Viewing the Details of a Bug

For the code to run, we must change Form\_Open:

```
IF (check_id() = FALSE) THEN
    message ("Unable to log on as " & system.user)
    ME.close
ELSE
    loadCombos
END IF
```

Now we can click on the combo-box to select the required bug ID, and use this to run another query in order to view the particular details. To do this, we need a text area (txtDetails) and some text boxes (txtCreated and txtRaisedby, txtPriority, txtApplication and txtVersion).

Double-click on cmbBugId, go into code edit mode, and edit the code so that it reads:

```
PUBLIC SUB cmbBugid_Click()

data.Exec ("SELECT l.priority,l.created, " &
    "l.details,l.release," &
    "s.title, c.username, a.name" &
    " from listing l,coder" &
    " c,status s,application a" &
    " where l.id=" & cmbBugid.Text &
    " AND l.status=s.id" &
    " AND l.raised_by_id=c.id"
    " and l.application_id = a.id")
```

```
txtDetails.Text = data.tmpRec!details
txtCreated.Text = data.tmpRec!created
txtRaisedby.Text = data.tmpRec!username
txtPriority.Text = data.tmpRec!priority
txtApplication.Text = data.tmpRec!name
txtVersion.Text = data.tmpRec!release
cmbStatus.Text = data.tmpRec!title
```

END

The next stage is to be able to log a new bug. We need to create a new form (frmAddBug), and we add an extra button to frmBugChecker—calling it btnAddBug and change the text to Add Bug. Don't forget to add a Close button before doing anything else. Next, add a text area (txtDetails), a text box (txtPriority) and a combo-box (cmbApplication). You also will need another button (btnSave):

```
PUBLIC SUB Form_Open()
```

```
    loadCombos
```

END

```
PRIVATE SUB loadCombos()
```

```
    data.loadCombo(cmbApplication,"select name from application")
```

END

Our third form will view all bugs. Create a new form (frmViewAll), and then go to frmBugManager, copy all of the elements and paste them into frmViewAll. You need to change the order of the

# LAYER 42™

- 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.

www.layer42.net

All prices include 100Mbps port, Firewall,  
24x7 Monitoring and DNS hosting

408-450-5740
2312 Walsh Ave., Santa Clara, CA 95051

objects so that the status combo-box is at the top of the form. Also add another text box (txtCoder). Finally, select the new Add Bug button, change the name to btnAccept and change the text to Accept.

For this form, we need some code to load cmbStatus first, because this will drive the others:

```
PUBLIC SUB Form_Open()
    loadCombos
    cmbStatus_Click
END

PRIVATE SUB loadCombos()
    data.loadCombo(cmbStatus,
        "select title from status")
END

PUBLIC SUB cmbStatus_Click()
    data.loadCombo(cmbBugid,"select l.id" &
        " from listing l, status s" &
        " where l.status = s.id" &
        " and s.title = '" & cmbStatus.Text & "'")
    cmbBugid_Click
    IF (cmbStatus.Text = "new") THEN
        btnAccept.Enabled = TRUE
    ELSE
        btnAccept.Enabled = FALSE
    END IF
END
```

Figure 5. The Final Application

```
PUBLIC SUB cmbBugid_Click()
    txtDetails.Text = ""
    txtCreated.Text = ""
    txtRaisedby.Text = ""
    txtCoder.Text = ""
    txtPriority.Text = ""
    txtApplication.Text = ""
    txtVersion.Text = ""
    IF ( cmbBugid.Text <> "" ) THEN
        data.Exec ("SELECT l.priority,l.created," &
            " l.details,l.release," &
            " s.title, r.username," &
            " c.username coder,a.name" &
            " from listing l,coder" &
            " r,status s,application a" &
            " left join coder c on l.coder_id=c.id" &
            " where l.id=" & cmbBugid.Text &
            " AND l.status=s.id" &
            " AND l.raised_by_id=r.id" &
            " and l.application_id = a.id")
        txtDetails.Text = data.tmpRec!details
        txtCreated.Text = data.tmpRec!created
        txtRaisedby.Text = data.tmpRec!username
        txtCoder.Text = data.tmpRec!coder
        txtPriority.Text = data.tmpRec!priority
        txtApplication.Text = data.tmpRec!name
        txtVersion.Text = data.tmpRec!release
    END IF
END

PUBLIC SUB btnAccept_Click()
    data.runSQL("update listing" &
        " set coder_id = " & data.coder_id(system.User) &
        ",status=2" &
        " where id = " & cmbBugid.Text)
    Form_Open
END
```

The last thing you need is for status change code in frmBugTracker:

```
PUBLIC SUB cmbStatus_Click()
    DIM version AS String
    version = txtVersion.Text
    IF (version="") THEN
        version = "Null"
    END IF

    data.runSQL(" update listing" &
        " set status = " & data.get_id("status","title", cmbStatus.Text) &
        " ,release = " & version &
        " where id = " & cmbBugid.Text)
END
```

This has been a very brief look at Gambas, but hopefully, it has shown just how easy it is to use to create a real working application. ■

**Resources for this article:** [www.linuxjournal.com/article/8951](http://www.linuxjournal.com/article/8951).

Mark Alexander Bain learned his trade working at Vodafone for nearly 20 years—using UNIX and Oracle. More recently, he has worked as a lecturer at a University in the UK and as a freelance writer. You can find his work at [www.markbain-writer.co.uk](http://www.markbain-writer.co.uk).

Compliance

Windows

RFID



Your best

EXPOSURE

Best Practices

Wireless

Linux

Success

C3 EXPO returns in 2006 as the premier IT event. It's the only computer industry event where corporate and channel buyers are exposed to the most innovative End-to-End Solutions driving the technology industry. All in one place. And at one time.



**Critical Linux**

The Linux Conference Track  
RedHat CEO Matthew Szulik  
Keynotes C3 EXPO



**Ian Murdock**, Free Standards Group CTO, Former Chairman and Chief Strategist of Progeny, and founder of the Debian Project, offers his vision on open source computing

to the world of IT

The Jacob K. Javits  
Convention Center, NYC  
June 27-29, 2006

Don't miss the New York area's only major event where you can explore practical methods, tools, and frameworks used to increase platform reliability, security, and asset life cycles using Linux platforms and tools - from these and many other industry leaders..



HA Bruno LLC  
210 East Route 4,  
Suite 304  
Paramus, NJ 07652

For information or to Register Online Now,  
visit [C3EXPO.COM/LJ](http://C3EXPO.COM/LJ) or call 201.543.5060 ext. 608

# How to Set Up and Use Tripwire

All about Tripwire and what it can do for you. MARCO FIORETTI

**Tripwire is an** intrusion detection system (IDS), which, constantly and automatically, keeps your critical system files and reports under control if they have been destroyed or modified by a cracker (or by mistake). It allows the system administrator to know immediately what was compromised and fix it.

The first time Tripwire is run it stores checksums, exact sizes and other data of all the selected files in a database. The successive runs check whether every file still matches the information in the database and report all changes. Tripwire initially was released in 1992. Today, several programs share this name, one is GPLed and two are proprietary. The rest of this article discusses only the GPL version 2.3.1.

## When Is the Right Moment to Start Using Tripwire?

IDS tools are particular beasts, and Tripwire is no exception. Even if you don't need to be an expert programmer to use this package, actually taking advantage of it requires some patience, attention and manual work.

First, using Tripwire is one of those cases in which blindly pressing Enter at every prompt really isn't a smart thing to do. Do yourself a favor and check at least the relevant parts of the good documentation provided with the Tripwire programs (more on this later).

Second, using Tripwire for real makes sense only if it is installed, fully configured and initialized at the very first boot after an installation from scratch, before ever connecting to the Internet or doing anything else. It takes only one attack to install a back door. All you would accomplish by installing and using Tripwire after such an event would be to guarantee that the back door remains just as open as the day a cracker installed it! Of course, even if you don't want to or can't re-install everything now, nothing prevents you from downloading the package anyway and becoming familiar with it.

Here is how to explain to Tripwire what's important to you. The Tripwire distribution includes several binaries, the corresponding man pages and two files that regulate the program's behavior, which we will call, for brevity, the Tripwire system files. The first one (`/etc/tripwire/twcfg.txt`), where several variables are defined, is for general configuration and even may be the same for all the computers on the same LAN. Its contents go from the location of the Tripwire database to instructions on minimizing the amount of time the passphrases are kept in memory or the number of redundant reports.

Other important parameters are the editor (the default is `vi`) for interactive usage and how reports should be sent by e-mail. The complete syntax and meaning of all possible variables are described in the `twconfig` man page.

The other system file (`/etc/tripwire/twpol.txt`) contains the policy that declares all the objects that must be monitored and what to do when one of them is lost or altered. Unlike the configuration file, the policy could (and almost certainly will) vary across the several computers on the same network. For example, the packages installed on a firewall will be different from those on a development workstation or an office laptop, even if the same GNU/Linux distribution is used.

The first thing to do to create a good Tripwire policy (and, in

general, have a less stressful sysadmin life) is to remove as many unneeded programs as possible before starting. Next, to make your usage of Tripwire as quick and effective as possible, your policy must cover everything you really need to monitor and nothing else. This includes, at least, all the system binary and library directories (that is, the contents of `/bin`, `/sbin`, `/usr/bin`, `/lib` and so on) and the corresponding configuration files in `/etc/`. The example `twpol.txt` files distributed with Tripwire contains *anything* that could be on a UNIX system, so it is guaranteed to complain about programs that you never installed or placed in a different location. This is an example of what you might see:

```
### Warning: File system error.
### Filename: /dev/cua0
### No such file or directory
### Continuing...
```

There is a safe and easy way, even if potentially long and boring, to remove such bogus warnings. Simply run the initial configuration procedure described below several times. Scan the report each time, and comment out the checks that generated false alarms until they all disappear. Of course, before starting, do what should be done before configuring any new package—that is, make a copy of the originals:

```
cp -p twcfg.txt twcfg.txt.orig
cp -p twpol.txt twpol.txt.orig
```

A Tripwire policy is a sequence of two kinds of rules. Normal ones define which properties of a file or directory tree must be checked, in this format:

```
object_name -> property_mask (rule attribute = value);
```

The `property_mask` specifies which properties must be examined or ignored. Attributes provide additional, rule-specific information like the rule severity or who should be informed by e-mail if that rule is violated. The other kind of rules are stop points, which tell Tripwire not to scan a particular file or directory. Tripwire also understands several directives for conditional interpretation of the policy, diagnostics and debugging. To know all the gory details, print out and study the `twpolicy` man page.

## Initial Configuration

After everything has been placed in the proper directories, either from a binary package or compiling the sources, the first action to take as root is to generate two robust—that is, hard to guess—passphrases. The first one (site passphrase) is used to encrypt and sign the Tripwire system files. The second one (local passphrase) is necessary to launch the Tripwire binaries.

Theoretically, the Tripwire distribution should include an `/etc/tripwire/twinstall.sh` script that should prompt the user for

passphrases and other information and then perform all the steps below. At the time of this writing, both the Tripwire 2.3.1 RPM package for Fedora Core 4 tested for this article and several on-line tutorials still say to use that script, but it just wasn't there after the installation. In any case, the utility that performs these tasks is twadmin. Because it has a complete man page, and it should be used anyway if you want to change keys after installation, we just show how it works. The actions described above are executed with the following commands:

```
twadmin --generate-keys --site-keyfile my_home_key
  ↳--site-passphrase 'Hello LJ readers'
twadmin --generate-keys --local-keyfile my_local_key
  ↳--local-passphrase 'Penguins are cool'
```

This leaves the two keys encoded in the my\_home\_key and, respectively, my\_local\_key files. Remember to copy these two names in the twcfg.txt file *before* running twadmin:

```
SITEKEYFILE      =/etc/tripwire/my_home_key
LOCALKEYFILE     =/etc/tripwire/my_local_key
```

Once the passphrases have been stored, the configuration file must be encrypted in this way:

```
twadmin --create-cfgfile --cfgfile twcfg.enc
  ↳--site-keyfile my_home_key twcfg.txt
Please enter your site passphrase:
Wrote configuration file: /etc/tripwire/twcfg.enc
```

The procedure to create a binary version of the policy is similar:

```
twadmin --create-polfile --cfgfile twcfg.enc --polfile
  ↳twpol.enc --site-keyfile my_home_key twpol.txt
```

The difference, with respect to the former command, is that now the encrypted configuration file must be passed to twadmin. The reason the two files must be encrypted is that Tripwire will discover if they are corrupted much more easily than if they were in plain-text format. In order to read such files directly, you need (besides the passphrases, obviously) the -print-cfgfile or --print-polfile options of twadmin.

### Database Creation

Once the passphrases and system files are all set, it's time to go into what the documentation calls Database Initialization Mode:

```
tripwire --init --cfgfile twcfg.enc --polfile tw.pol
  ↳--site-keyfile my_home_key --local-keyfile my_local_key
```

By default, the result is stored in /var/lib/tripwire/YOURHOSTNAME.twd. Path and name can be changed in twcfg.txt or given as a command-line option. Eventually, if everything goes fine, you'll be greeted by this message:

```
Wrote database file: /var/lib/tripwire/YOURHOSTNAME.twd
The database was successfully generated
```

### Periodic Checks

As soon as encrypted system files, passphrases and a complete snapshot of your system are available, Tripwire finally can do the only thing we really care about—that is, to check the integrity of our computers periodically. This is normally accomplished by running the program as a




Figure 1. Sample Tripwire Report

cron job with this switch:

```
tripwire --check
```

Note that, just to allow secure, automatic usage, the program doesn't need passphrases when launched in this way. Consequently, there is no need to write them in plain text anywhere. The integrity report is printed both to STDOUT (so it can be e-mailed to the system administrator) and saved in the location specified by the REPORTFILE variable in the configuration file. How often this operation should be performed depends on how critical the system is and how often it is exposed to external attacks. Although a corporate firewall should be checked daily, a weekly check may be enough for a department print server behind it or a regular desktop.

Figure 1 shows an example of what a Tripwire report looks like. It tells you, for every rule defined in the policy, which of the corresponding files were added, changed or modified. Command-line options are available to check only specific sections of the policy file, or just some files. This could be useful, for example, when nothing was modified in the system, but there is the suspicion that some particular disc or partition was damaged.

The integrity checking procedure also can be interactive. Adding the -interactive switch causes Tripwire to open an editor, after the check, to allow the user to declare which files should be permanently updated in the Tripwire database. This is a manual alternative to the update mode described below.

### Update Policies

Immediately after any system change, be it due to installation, update or removal of software or configuration files, it is mandatory to update the plain-text policy file and regenerate the binary database. Any successive Tripwire check would be meaningless otherwise. Therefore, run this command whenever it's necessary:

```
tripwire -update-policy -twrfile
  ↳a_previous_integrity_report.twr
```

Because it is so critical, this operation requires both your local and site passphrases. When launched in this way, Tripwire detects as violations any changes that happened after the specified integrity check. In such a case, an actual update of the policy, ignoring such violations, is possible only if the user explicitly tells the program to run in low security mode. The corresponding option is `-Z low` and is explained in detail in the Tripwire man page.

### Miscellaneous Tips and Tricks

Reading the `twfiles` and `twintro` man pages, which contain short and up-to-date overviews of all the files and programs that compose the Tripwire suite, is highly recommended before starting the installation. The actual Tripwire binary, if called with the `-help` option, lists all the available options. Like many FOSS programs, all the utilities of this package accept both short and long forms of their command-line options.

For example, `tripwire -check` also can be written as `tripwire -m c`. The second form is faster when one already knows Tripwire and has to use it interactively, but the explicit command is recommended in scripts, for documentation or didactical purposes. The `-v` option puts any Tripwire command in verbose mode. Common wisdom also suggests that both the binary and text versions of the Tripwire system files be stored on a separate computer, write-protected floppy disk or USB drive.

Remember that one of the first things a determined cracker will do is to replace just *those* files with her own copies, to hide any trace of attack. The periodical reports placed by Tripwire in `/var/lib/tripwire` are in binary, optionally signed format. Consequently, they can't be read straight from the prompt, and they also can't even be processed directly by a shell script for automatic comparison or other purposes. The solution is to use the `twprint` command, which comes with its own complete man page, as in this example (note that you must pass the binary configuration file for it to work):

```
twprint --print-report --cfgfile twcfg.enc --twrfile
➔ /var/lib/tripwire/report/my_tripwire_report
```

The digital signatures of each binary file can be checked directly with the `siggen` utility, which also has its own man page:

```
/usr/sbin/siggen /etc/tripwire/twcfg.enc
-----
Signatures for file: /etc/tripwire/twcfg.enc

CRC32          Dmjk1z
MD5            DTn311w6Wx3+7TXv75HPjA
SHA            D5N1Pv4biCnd14igf/anGM3pvVH
HAVAL          BEJmfzpcA/Txq5nf9kgsVb
```

### Credits and a Glimpse of What's Ahead

The Open Source Tripwire Project had been quiescent for some time. Luckily, just a few days before the deadline of this article, version 2.4.0.1 was released on SourceForge, and it is the one you'll likely find packaged for your distribution by the time you read this. Besides the source tarball, it is also possible to download x86 static binaries built on a Gentoo 2005 distribution. There are no remarkable changes in functionality, so everything explained in this article should still apply as is. The other good news is that this is the first release in which the old build system has been replaced by a standard `autoconf/configure` environment. Unfortunately, due to some `gcc 4` compatibility problems on Fedora Core 4, it wasn't possible to test this version in time. However, as soon as this porting is completed, it should be much easier to add new features and package Open Source Tripwire for all modern GNU/Linux distributions. You're welcome to join the effort and report bugs on the developers mailing list (see the on-line Resources). Thanks to Paul Herman and Ron Forrester for releasing this new version and the time they spent to answer my questions. ■

**Resources for this article:** [www.linuxjournal.com/article/8950](http://www.linuxjournal.com/article/8950).

---

Marco Fioretti is a hardware systems engineer interested in free software both as an EDA platform and, as the current leader of the RULE Project, as an efficient desktop. Marco lives with his family in Rome, Italy.

## TRIPWIRE COMPETITORS

Tripwire certainly isn't the only IDS solution for GNU/Linux systems. Its more popular competitors are briefly described below. Other similar packages in the same category are listed on a Purdue University Web page (see Resources).

**AIDE:** the Advanced Intrusion Detection Environment is proposed as a free software replacement for everything Tripwire does and then some. The integrity database is defined by writing regular expressions in the configuration file. Besides MD5, several other algorithms can be added to verify the file integrity. Today, AIDE already supports sha1, `tormd160`, `tiger`, `haval` and some others.

**Tripwire for Servers:** this is a proprietary product by Tripwire, Inc., the company created by the original Tripwire developers. It is targeted at centralized monitoring of groups of Linux, UNIX or Windows servers.

**Tripwire Enterprise:** the flagship product of Tripwire, Inc., is for automated monitoring of mixed networks of up to thousands of servers, desktops, directory servers and network devices. Starting with version 5.2, Tripwire Enterprise includes centralized management through a Web interface, as well as customizable reporting and control.

# Why is LPI the World's #1 IT certification?

---

## Stable.

All Linux Professional Institute certification programs are created with extensive community and enterprise input, rigorous psychometric scrutiny, and professional implementation procedures

LPI strives to remain an independent and unbiased certification body. As a result, LPI is supported by a broad range of enterprise customers, governments, testing centers, book publishers, courseware suppliers, and educational institutions from around the globe.

## Innovative.

LPI's programs follow Linux Standard Base (LSB) specifications, so people who earn our certifications are qualified to work on all major distributions. Through our strong grass-roots base and corporate support both inside and outside the world of open-source, LPI goes beyond "vendor-neutral" to truly address community and enterprise needs.

We are the first IT certification to obtain professional accreditation and advocate the adoption of Open Source Standards through our work with such organizations as the Free Standards Group. We also committed to the development of innovative open source software tools which will improve and streamline our test development processes.

## Growing.

We have delivered over 100,000 exams and 30,000 certifications around the world. Our exams are available in several languages, at more than 7,000 locations, in more than 100 countries. You can take LPI exams where and when you want.

LPI serves Linux professionals and the IT industry. We have attracted broad support from prominent members of the Linux community and enterprise corporations, as demonstrated by our Strategic Advisory Council and sponsors. In addition, we are establishing a Technical Advisory Council to ensure we address the future needs of the industry. Our vendor-independent stance ensures we focus solely on the skills and knowledge necessary for an IT professional rather than on the promotion of a specific software vendor or distribution.



**Linux  
Professional  
Institute**

For more information,  
Please contact us at  
[Info@lpi.org](mailto:Info@lpi.org) or visit us at  
[www.lpi.org](http://www.lpi.org).

---



# The World Is a libferris Filesystem

With libferris, the boundary of your filesystem extends to include PostgreSQL, XML, db4, RDF, the X Window System, Evolution and much more. BEN MARTIN

The libferris virtual filesystem always has sought to push the boundaries of what a filesystem should do in terms of what can be mounted and what metadata is available for files. During the past five years, it has expanded its capabilities from mounting more traditional things, such as tar.gz, SSH, digital cameras and IPC primitives, to being able to mount various Indexed Sequential Access Mechanism (ISAM) files, including db4, tdb, edb, eet and gdbm; various relational databases, including odbc, MySQL and PostgreSQL; various servers, such as HTTP, FTP, LDAP, Evolution and RDF graphs; as well as XML files and Sleepycat's dbXML.

Recently, support for indexing filesystem data using any combination of Lucene, ODBC, TSearch2, xapian, LDAP, PostgreSQL and Web search has been added with the ability to query these back ends for matching files. Matches naturally are presented as a virtual filesystem. Details of using the index and search capabilities of libferris appeared in the February 2005 issue of *Linux Journal* in my article "Filesystem Indexing with libferris". I should mention that anything you see mounted as a filesystem in this article can be indexed and searched for as described in that past article on searching.

You can access your libferris virtual filesystem either by native libferris clients or by exporting libferris through Samba.

The two primary abstractions in libferris are the Context and the Extended Attribute (EA). A Context can be thought of as a superclass of a file or directory. In libferris, there is less of a distinction between a file and a directory with the ability for a file to behave like a directory if it is treated like one. For example, if you try to read a tar.gz file as a directory, libferris automatically mounts the archive as a filesystem and lists the contents of the archive as a virtual filesystem.

The EA interface can be thought of as a similar concept to the Linux kernel's EA interface. That is, arbitrary key-value data is attached to files and directories. This EA concept was extended early on in libferris to allow the value for an attribute to be derived from the content of a file. This means simple things like width and height of an image or video file become first-class metadata citizens along with a file's size and modification time. The limits on what metadata is available extend far beyond image metadata to include XMP, EXIF, music ID tags, Annodex media, geospatial tags, RPM metadata, SELinux integration, partially ordered emblem categories and arbitrary personal RDF stores of metadata.

Having all metadata available through a single interface allows libferris to provide filtering and sorting capabilities on any of that metadata. As such, you can sort a directory by any metadata just as easily as you would use `ls -Sh` to sort by file size. Sorting on multiple metadata values is also supported in libferris; you can sort your files easily by MIME type, then image width, then modification time—with all three pieces of metadata contributing to the final directory ordering. Any libferris virtual filesystem can have filtering and sorting applied to it to obtain a new libferris virtual filesystem.

You can store EA values into a personal RDF store—for example, when you write an image width to an extended attribute. When you subsequently read the image width, you get the value you just wrote to the EA. This extends naturally to other situations, such as when you

change the x or y EA for a window, which should move the window.

Allowing EA to be stored in a personal RDF file lets you add metadata to any libferris object, even those for which you have only read access. For example, you can attach emblems or comments to the Linux Kongress Web site just as you would a normal file.

An interesting EA for all files is the content EA, which is equivalent to the file's byte contents. Exposing the file itself through the EA interface means that any information about a file can be obtained via the same interface.

libferris is written in C++ and provides a standard IOSTream interface to both Contexts and EA. Many standard file utilities have been rewritten to take advantage of libferris features. These clients include `ls`, `cp`, `mv`, `rm`, `mkdir`, `cat`, `find`, `touch`, IO redirection and more.

## Filesystem Interaction

As we explore these filesystems, I use the `ferrisls` command, which mimics the `coreutils ls(1)` command. As well as the `-l` long listing option, I use the `-0` (zero) recommended-`ea` option of `ferrisls`. This operates in much the same way as `-l`, though it asks the filesystem itself which EAs are most interesting for the user to see. I assume a shell alias of `fls=ferrisls` in the code examples.

I start by showing interaction with the standard kernel-based filesystems and some of the EA possibilities. Along with the recom-

Listing 1.

### A Long Listing of a Directory with Explicit Metadata

```
$ fls -l \
--show-ea=size-human-readable,width,height,name
4.5k 48 46 emacs.png
1.9k 48 48 gnome-warning.png
3.2k 48 48 gnome-xterm.png
2.5k 48 48 gtkvim.png
```

Listing 2.

### Asking libferris itself to determine which EAs are of interest for the current directory and producing an XML document as output.

```
$ fls -0 --xml
<ferrisls>
<ferrisls url="file:///tmp/lj" name="lj" >
<context size-human-readable="4.5k"
  protection-ls="-rw-r-----"
  mtime-display="05 Dec 4 23:39"
  name="emacs.png" width="48" height="46" />
...
</ferrisls>
</ferrisls>
```

mended-ea option, ferris supports the --xml option to produce an XML document as output. This provides information as to what EA each value belongs and provides one possibility to drive Web interfaces using libferris.

As mentioned previously, if you are sorting a directory on an EA that does not provide a complete ordering, you can chain together sorting predicates. For example, in Listing 3, I have sorted the output based on the numeric EA height and then used a version string sort on the name EA. A version sort is similar to the ls(1) -v option, which in Listing 3 has placed foo20.png after foo3.png. Such sorting is very useful when sorting by file type or MIME major type followed by name.

### Listing 3. Sorting Your Output

```
$ fls --show-ea=width,height,size,name \
--ferris-sort='(:#:height)(:V:name)'
```

48	48	1968	gnome-warning.png
48	48	3253	gnome-xterm.png
48	48	2550	gtkvim.png
48	46	4589	emacs.png
48	46	4589	foo3.png
48	46	4589	foo20.png

The two concepts of files forming a tree and files having key-value pairs attached to them are similar to the structure of XML. With libferris, you can poke inside XML documents as though they were just another filesystem. For example, see Listing 4.

### Listing 4. Initial Exploration of XML as a Filesystem

```
$ cat example.xml
<root>
  <file1 size="200" />
  <file2 interesting="yes" />
  <file3>filesystems rock
</file3>
</root>

$ fls -0 ./example.xml/root
file1
file2
file3

$ fls -d --show-ea=name,interesting \
./example.xml/root/file2
file2 yes

$ fcat example.xml/root/file3
filesystems rock
```

By interacting with your filesystem, you can cause updates on the underlying XML document as well. The ferris-redirect client exists to allow shell-like redirection into libferris files. The -T or --trunc option truncates an existing file before writing stdin into it. This is much like the >| shell option. As you can see from the interaction in Listing 5, we have changed the structure of the example.xml document significantly through filesystem interaction.

## 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 #
AML www.amltd.com	31	MICROWAY, INC. www.microway.com	C4, 71
APPRO HPC SOLUTIONS appro.com	C2	MIKRO TIK www.routerboard.com	7
ASA COMPUTERS www.asacomputers.com	55, 59	MONARCH COMPUTERS www.monarchcomputer.com	20, 21
C3 EXPO www.C3EXPO.COM	75	OPEN SOURCE STORAGE www.opensourcestorage.com	13
CARINET www.cari.net	67	PENGUIN COMPUTING www.penguincomputing.com	29
CIARA TECHNOLOGY www.ciaa-tech.com	8, 9, 16, 17	POLYWELL COMPUTERS, INC. www.polywell.com	37
CORAID, INC. www.coraid.com	39	THE PORTLAND GROUP www.pggroup.com	35
COYOTE POINT www.coyotepoint.com	3	QSOL.COM www.qsol.com	53
CYCLADES CORPORATION www.cyclades.com	1	RACKABLE SYSTEMS www.rackable.com	57
EMAC, INC. www.emacinc.com	36	RACKSPACE MANAGED HOSTING www.rackspace.com	C3
EMPERORLINUX www.emperorlinux.com	15	R CUBED TECHNOLOGIES www.rcubedtech.com	61
FAIRCOM CORPORATION www.faircom.com	45	SBE, INC. www.sbei.com	23
FREE SOFTWARE FOUNDATION member.fsf.org	95	SERVERS DIRECT www.serversdirect.com	43
GENSTOR SYSTEMS INC. www.genstor.com	49	SUGARCRM, INC. www.sugarcrm.com	69
HURRICANE ELECTRIC www.he.net	27	SUPERMICRO www.supermicro.com	47
INTEGRITY LINUX SYSTEMS www.integritylinux.com	51	TECHNOLOGIC SYSTEMS www.embeddeddx86.com	83
INTEL www.intel.com/software/products	11	THINKMATE www.thinkmate.com	19
IRON SYSTEMS www.ironsystems.com	41	TUX MAGAZINE www.tuxmagazine.com	87
LAYER 42 NETWORKS www.layer42.net	73	TYAN COMPUTER USA www.tyan.com	5
LINUX JOURNAL www.linuxjournal.com	12, 85	ZERVEX www.zervex.com	6
LPI www.lpi.org	79	ZT GROUP INTERNATIONAL www.ztgroup.com	33
MBX www.mbx.com	25		

Listing 5.

## Changing an XML File through Its Filesystem

```
$ echo "VIRTUAL filesystems rock more" | \
  ferris-redirect -T ./example.xml/root/file3

$ echo "a new way" | \
  ferris-redirect ./example.xml/root/file4

$ ferrisrm ./example.xml/root/file2

$ ftouch ./example.xml/root/touched

$ cat example.xml
<?xml version="1.0" encoding="UTF-8"
  standalone="no" ?>
<root>
  <file1 size="200"/>

  <file3>VIRTUAL filesystems rock more
</file3>
  <file4>a new way
</file4>

  <touched/>

</root>
```

Listing 6.

## OpenOffice.org Documents Are Filesystems Too

```
$ fls -lh show-ea=size,name,content \
~/sample-oo-writer.odt/content.xml/ \
office:document-content/office:body/office:text
0      office:forms
18     text:p Paragraph number 1
0      text:p-1
116   text:p-2 This is the second paragraph ...
0      text:p-3
39    text:p-4 And in summary, this is really...
0      text:p-5
0      text:sequence-decls
```

As many modern word-processing documents are XML inside a compressed container, libferris allows you to drill down into the office document as though it were a filesystem. In Listing 6, I am listing a simple OpenOffice.org Writer document as a filesystem.

A Xerces-C Document Object Model (DOM) can be obtained for any libferris filesystem, just as a Xerces-C DOM can be mounted as a libferris filesystem. Creation of a DOM for a filesystem is evaluated lazily, so you can get a DOM for file:// and only the parts of the DOM that are required are ever created.

The ability to convert any libferris filesystem into a DOM allows you to apply XSLT to your filesystems easily. The example C++ code in Listing 7 applies a stylesheet to a mounted OpenOffice.org document.

Recently, support for mounting applications, such as Firefox, Evolution and the X Window System, was added to libferris.

The evolution:// filesystem allows you to mount your Evolution mail client. Support currently extends to your mail folders and contacts. Using this filesystem, it is no longer necessary to save attachments to temporary files to access them from ferris-aware systems.

Mounting your X Window System is done via the xwin:// filesystem.

Listing 7.

## C++ Code Fragment Applying an XSLT to a Filesystem

```
fh_context c = Resolve( "~/example.odt/content.xml/"
  "office:document-content/office:body/office:text");
DOMDocument* theDOM = Factory::makeDOM( c );
...
// should use XercesDOMWrapperParsedSource
XalanTransformer theXalanTransformer;
theXalanTransformer.transform(
  theDOM, "~/my-oo.xsl", cout );
```

Listing 8.

## Mounting Evolution and the X Window System

```
$ fls evolution://localhost
contacts mail
$ fls -0 evolution://localhost/contacts/system/
...
witme-ferris  witme-ferris@lists.sourceforge.net
...
$ fls -0 xwin://localhost/clipboard
0      #include <Ferris/Ferris.hh>
1      Let the cricket stick to its hearth
2      ...
```

Listing 9.

## PostgreSQL as a Filesystem

```
$ psql
# create database tmp;
# \c tmp
# create table foo ( message varchar(100) not null,
                        id int primary key );
# insert into foo values ( 'doki doki', 1 );
# \q

$ fls -0 pg://localhost/tmp/foo
doki doki 1      1      id

$ fcat pg://localhost/tmp/foo/1
<context id="1" message="doki doki" />

$ echo "waku waku" | ferris-redirect \
  -T --ea=message pg://localhost/tmp/foo/1

$ fls -0 pg://localhost/tmp/foo
waku waku 1      1      id

$ gfcreate pg://localhost/tmp/foo
# See the gfcreate-tuple figure

$ fls -0 pg://localhost/tmp/foo
utsukushii 2      2      id

$ psql tmp;
# select * from foo;
      message          | id
-----+-----
waku waku              | 1
utsukushii             | 2
```




Figure 1. Creating a New Tuple in PostgreSQL through the Filesystem

This gives access to your window objects as well as lets you mount Klipper on KDE desktops. For Klipper, you can `ls`, `cat` and `cp` your past clipboard interactions easily, and overwriting the top clipboard element is effectively a clipboard copy. The window mounting lets you see where your windows are in terms of x,y offsets as well as other interesting data. Listing 8 shows a sample session of mounting my Evolution mail client and the X Window System.

Mounting databases allows you to explore the database server, its databases and their tables and views. Shown in Listing 9, I create a database, populate it and interact with it as a virtual filesystem. The final command using the `--xml` option for `ferrisls` exports each tuple in XML format.

Instead of embedding the user name and password in the URL, `libferris` elects to store this information in configuration files. This is a trade-off when the risk of accidentally copying and pasting a URL with embedded user credentials is minimized at the expense of having a central store of available credentials and mappings for where to use each credential. For many common URLs, inline authentication information is also supported.

The invocation of `gfcreate` shown in Listing 9 is captured in Figure 1.

Listing 10.

What types of things can I create for a PostgreSQL filesystem?

```
$ fcreate -l pg://localhost/tmp
listing types that can be created
for context: pg://localhost/tmp
queryview
table

$ fcreate -l pg://localhost/tmp/foo
listing types that can be created
for context: pg://localhost/tmp/foo
tuple
```

A `libferris` filesystem can nominate which objects it is happy to have created on it. You can see this list by using the `fcreate` or `gfcreate` tools in the `ferriscreate` package. A large list of possibilities will be displayed for an `fcreate -l /tmp`, for example. For a PostgreSQL database, you can create only a small number of new object types, as shown in Listing 10. I'll use `fcreate` in a moment to create a new empty `db4` file to show how filesystem monitoring is virtualized in `libferris`.

Many changes made to a `libferris` filesystem are reflected instantly in other `libferris` applications. Many kernel-level interfaces let applications know when a kernel filesystem changes—for example, `inotify` and `dnotify`. `libferris` extends this to allow clients to know when a virtual filesystem has changed. For example, when you update an element inside of an XML file, `inotify` tells you only that the XML file has changed. With `libferris`, you can see exactly which part of the XML file was modified by other `libferris` applications.

Listing 11 demonstrates the filesystem monitoring support. In the example, I use the `--monitor-all` option of `ferrisls`. This makes `ferrisls` operate like a `tail -f` for your given URL; any creation, deletion or

Listing 11.

Output of One Virtual Console

```
$ fcreate --create-type=db4 --rdn=raw.db .
$ fls --monitor-all -0 ./raw.db
Created new1
Changed c:0x8321f88 /tmp/ljdb/raw.db
Changed c:0x8321f88 /tmp/ljdb/raw.db
Deleted new1
Created redirected-output
Changed c:0x8321f88 /tmp/ljdb/raw.db
```

**NEW!**

TS-7300 High Security Linux FPGA Computer



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

**200 MHz CPU**

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

**Design your solution with one of our engineers**

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

**See our website for options, peripherals and x86 SBCs**



We use our stuff.  
visit our TS-7200 powered website at  
[www.embeddedARM.com](http://www.embeddedARM.com)  
(480) 837-5200

Listing 12.

**Output of Another Virtual Console**

```
$ ftouch ./raw.db/new1

$ ferrisrm -v ./raw.db/new1
removing ./raw.db/new1

$ echo "hello" | \
  ferris-redirect -T ./raw.db/redirected-output

$ fcat ./raw.db/redirected-output
hello
```

Listing 13.

**Setting Up Xine to Play Annodex Files**

```
$ cat xine.desktop
[Desktop Entry]
Name=xine
Comment=Video Player
Exec=xine
MimeType=video/mpeg;...
Icon=~/.icons/xine.png
Terminal=0
Type=Application
$ ferris-import-desktop-file xine.desktop
$ ferris-set-file-action-for-type -v -a xine \
  /tmp/Wombats.anx

# Lets view the video.
$ alias fv="ferris-file-action -v"
$ fv /tmp/Wombats.anx
```

interesting filesystem activity is shown on the console. In another terminal, Listing 12, I'm creating, deleting and writing to "files" inside a Berkeley db4 file. ferrisls happily reports what is happening to these virtual files.

Many operations performed with libferris are also stored for possible future use. This includes the types of files you recently created (png, jpeg, db, tuple and so on), which files you recently edited and viewed and more. All of this is kept only for your personal use and never sent anywhere. Storage of metadata on files you view and edit is called remembrance in libferris. Only view and edit actions invoked through libferris are currently remembered. Listing 13 shows how I set up Xine to be executed as the default view operation on Annodex media files.

Now we can explore what libferris knows about our past operations. By default, remembered operations are grouped by operation type then media type. The recommended EA for the final directories in the tree are the filename and the time it was last viewed or edited. This history virtual filesystem shown in Listing 14 shows only a set amount of the most recent operations so as not to become too large.

For each file, you also can bring up the complete list of view and edit times. This uses what libferris calls a branch filesystem. A branch filesystem best can be described as an entire personal filesystem attached to a file. Branch filesystems are accessed using the branches:// handler; all other URL handlers appear as direct children of branches://.

Listing 14.

**Showing Recent View Operations**

```
$ fls remembrance://
history
$ fls remembrance://history
edit view
$ fls remembrance://history/view
video
$ fls -0h remembrance://history/view/video
/tmp/Wombats.anx 05 Dec 6 21:34
```

Listing 15.

**Branch Filesystems: a Filesystem about a File**

```
$ fls branches://file/tmp/Wombats.anx
branchfs-attributes  branchfs-medallions
branchfs-remembrance branchfs-extents
branchfs-parents    branchfs-signatures
$ fls -0 branches://file/tmp/\
  Wombats.anx/branchfs-remembrance/view
10.7M -rw-rw---- 05 Dec 6 21:34 ... 05 Dec 6 21:35
10.7M -rw-rw---- 05 Dec 6 21:34 ... 05 Dec 6 21:39
...
$ fls --xml \
  branches://file/tmp/Wombats.anx/branchfs-extents
<ferrisls>
<ferrisls
  url="branches://.../branchfs-extents"
  name="branchfs-extents" >
<context name="0"
  start-block="14245376"
  end-block="14267375"
  start-address="0"
  end-address="21999" />
</ferrisls>
</ferrisls>
```

In Listing 15, I take a look at what branches are available for my media file and explore the remembrance view filesystem. Then, out of curiosity, I take a look into the extents branch and see that the kernel's XFS filesystem has placed the whole media file in a single contiguous extent on disk.

To see if a file has a valid digital signature, you simply can read the has-valid-signature EA on the file. The signatures branch filesystem allows much more detail to be exposed about the signature. The branchfs-attributes filesystem exposes all EAs for a file as a filesystem. Sometimes it is more convenient to access an EA as though it were a file.

**Future Directions**

In the future, libferris will continue to support mounting more things and obtaining more metadata where it can. A module for FUSE is planned to supplement the current Samba support. ■

**Resources for this article:** [www.linuxjournal.com/article/8947](http://www.linuxjournal.com/article/8947).

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-filesystem interaction.

**Do you take**

***"the computer doesn't do that"***

**as a personal challenge?**

**So do we.**

**LINUX**  
**JOURNAL**™

Since 1994: The Original Monthly Magazine of the Linux Community

**Subscribe today at [www.linuxjournal.com](http://www.linuxjournal.com)**

# USB Pendrives and Distributions for Them

A look at three distributions you can use to boot Linux from a USB pendrive.

JUAN MARCELO RODRIGUEZ

A **pendrive** is a USB storage device. You plug it in to a USB port, and if the pendrive is compatible with your operating system, it should look exactly like another disk on your system. These days, it is easy to find pendrives with 1GB of storage.

It so happens that there has been an explosion of bootable live CD versions of Linux. Both commercial and noncommercial Linux distributions are providing live CDs (including Linspire, SUSE, Ubuntu, Kubuntu, Knoppix and Mepis, to name only a few—there are many more).

Imagine a mixture of both concepts—a USB storage device and a live CD version of Linux. You can pack a lot of features of a GNU/Linux live CD into 1GB. The USB pendrive has the advantage of being writable, which the live CD lacks. So, you can boot Linux from a pendrive and store data on it too. The end result is that, as long as you can find a machine that will boot from a pendrive, you have a fully portable version of Linux that carries your applications, settings and data.

## The Choices

I cover three LiveUSB distributions in this article: SLAX, Damn Small Linux (DSL) and Flash Linux. Each one has different window managers and different apps.

SLAX works with tmpfs and Unification fs (UFS), which give it some nice advantages. SLAX is based on Slackware Linux with the 2.6 Linux kernel.

DSL is a little distribution of 50MB. DSL configures Fluxbox very nicely. Some of the apps included are Mozilla Firefox, the Slypheed mail client, xmms, text editors, graphics viewers and more. It includes a 2.4 Linux kernel with good hardware detection, but it doesn't have the big apps other distributions have, such as The GIMP. It is a compact distribution with a script to install it to LiveUSB.

Finally, Flash Linux is a solid distribution that uses the 2.6 kernel and the fast JFFS2 filesystem. It has good speed, both as a live CD and LiveUSB, and it includes large applications, such as The GIMP and OpenOffice.org. It uses grub, bootsplash, framebuffer and GNOME, and is based on Gentoo.

## Boot from USB

The biggest challenge in using a USB pendrive for your Linux distribution is booting the pendrive. Old motherboards do not support the ability to boot from USB hardware, so you may need to use a floppy disk to boot your USB-based distribution. Newer motherboards let you boot drives usually referred to as USBHDD, USBZIP, USB-FDD and others, such as USB-CDROM.

The first step to using a pendrive is to delete the original pendrive partitions, if there are any. Then, add a FAT16 partition, and

format it with mkdosfs. I used cfdisk to do the work, but you can use fdisk too.

Check your dmesg log when you plug in the device to see if it is working:

```
dmesg | tail
```

You should see a message similar to the following:

```
sda: assuming drive cache: write through
sda: sda1
```

Format the partition you created with the following line:

```
mkdosfs -F 16 /dev/sda1
```

(Change sda1 to whatever partition is appropriate for your system.)

Unplug the hardware, and plug it in again. You now are ready to install the distribution.

## Damn Small Linux

Go to the DSL Web site (see the on-line Resources) and download the ISO image file for the current version of DSL, and burn the ISO to a CD or DVD. Boot from this CD or DVD. The boot starts with a welcome screen, like most live distributions.

DSL looks for hardware, and then it installs and configures it. Depending on your machine, it will bring up an X server running Fluxbox in less than two minutes.

After booting from the DSL live CD, right-click on the Fluxbox desktop to open the Fluxbox menu. Go to Apps→Tools→Install to install it to your USB pendrive. Here, you have two options for installing the distribution: install to USBHDD or USBZIP hardware. DSL will ask about the location of the pendrive, and it also asks if you want to install DSL from the live CD, from a file or from the Web.

I suggest you use your broadband connection to download the files. In fact, if you have a router that supports DHCP, DSL should recognize your Ethernet card and have no problems accessing the Internet at boot time. DSL supports PPPoE too, if your Internet connection requires it.

I missed the features of the 2.6 kernel (the next release of DSL should support 2.6), but it's still a good little distribution. I think DSL is fine as it is, but if you need a big office suite, you should use SLAX. Resources that you *must* read if you use DSL are the Wiki and the complete DSL forums. You will find many tips and tricks with plenty of information that will be helpful if you run into problems.

## Floppy-Based Boot Process

If your machine doesn't allow you to boot from a USB pendrive, you can boot DSL from a floppy. Download the file `bootfloppy-usb.img` from the DSL site, and copy the image to a floppy disk with `dd`:

```
dd if=bootfloppy.img of=/dev/fd0
```

Modify your computer BIOS to boot from the floppy first, and then boot the floppy image file of DSL. This boot image will launch the USB version of DSL. This process works with just about any distribution that offers a floppy boot image for booting USB pendrives.

## SLAX

The SLAX site says, "SLAX is a fast and beautiful Linux operating system, which fits on small (3.14") CD-ROM disc. It runs directly from the CD (or USB) without installing. The live CD described here is based on the Slackware Linux distribution and uses the Unification File System (also known as `unionfs`), allowing a read-only filesystem to behave as a writable one, saving all changes to memory." Fortunately, when you use a pendrive, you don't have to worry about emulating write operations because, unlike a CD, the pendrive memory is writable.

You can use UFS to merge storage from several sources, including network storage, into one local directory. This makes UFS a good solution for diskless workstations, because it makes it easy to keep your home directory on a network storage device.

SLAX is a modular distribution, so you can add features as you need them. It lets you configure your installation for many different purposes. You might be able to watch a DVD, use QEMU, burn CDs and DVDs, run firewalls, antivirus apps and much, much more. Check the list available on the project's site (see Resources) to find out about modules that add new features to SLAX.

## SLAX Installation

To install SLAX, get the latest version from the Web site (see Resources). SLAX has many versions of the same distributions, with certain differences in apps and size. Select among Frodo, Standard, Popcorn or KillBill editions. I used the `slax-5.0.7b.iso` standard edition of 200MB with KDE.

Mount the ISO image file of SLAX using the loopback device. In my case, I called the mount directory `slax`. Here is the command I used:

```
mount -o loop slax-5.0.7b.iso slax/
```

As before, format the USB pendrive to use FAT16:

```
mkdosfs -F 16 /dev/sda1
```

(Change `sda1` to whatever partition is appropriate for your system.)

After you have a bootable and formatted FAT16 partition in the pendrive, mount it:

```
mount -t vfat /dev/sda1 /mnt/usb/
```

Copy all the files from the directory `slax/`, where you mounted the ISO of SLAX, to the mounted pendrive:

```
cp -rav slax/* /mnt/usb/
```

Synchronize the data:

```
sync
```

**Free Subscriptions!**

Dear Bill,

It's over between us.  
I've found someone new.  
Someone I can depend on.  
Someone who is fun for  
a change. Thought you might  
like to see his picture.

-Sandy

**TUX**

The first and only magazine for the new Linux user.  
Your digital subscription is absolutely free!  
**Sign up today** at [www.tuxmagazine.com/subscribe](http://www.tuxmagazine.com/subscribe)

And go to the pendrive location (/mnt/usb):

```
cd /mnt/usb/
```

Now, copy the files vmlinuz and initrd.gz to the root directory, where you mounted the pendrive, in our case from the directory /mnt/usb/, and do:

```
cp boot/vmlinuz .
cp boot/initrd.gz .
```

Then, edit the file called isolinux.cfg:

```
pico isolinux.cfg
```

Remove every string called boot/ before vmlinuz and initrd.gz. Then, rename it to syslinux.cfg to use syslinux with the device:

```
mv isolinux.cfg syslinux.cfg
```

Finally, install and update MBR with LILO or GRUB:

```
lilo -M /dev/sda
```

And, use syslinux to finish the process:

```
syslinux -s /dev/sda1
```

SLAX is installed—enjoy it. Umount the pendrive and reboot. Change your BIOS to boot from the USB pendrive, and reboot again. You may need to use LILO or GRUB to update or install the master boot record on the pendrive.

SLAX has KDE, Fluxbox, K3b, Media Player, a Web browser, mail, office suite, Kopete and many other applications. You can find a complete list on the SLAX Web site (see Resources).

SLAX doesn't have the speed of DSL, but has the 2.6.15 kernel, excellent network support, the parted application (partition editor) and more. It's by far a more complete distribution than DSL, but you pay for it in size.

## Flash Linux

The Flash Linux distribution is based on Gentoo Linux. Get the Flash Linux ISO image file from the Web site (see Resources), and burn it to CD. Then, boot from the CD in order to install the LiveUSB version in the pendrive. Download the three parts of the ISO from [sourceforge.net/project/showfiles.php?group\\_id=124770](http://sourceforge.net/project/showfiles.php?group_id=124770). Currently, the three parts are flashlinux-0.3.4-RC2.iso-part1, flashlinux-0.3.4-RC2.iso-part2 and flashlinux-0.3.4-RC2.iso-part3. After you download these files, put them together:

```
cat flashlinux-0.3.4-RC2.iso-part1 flashlinux-0.3.4-RC2.iso-part2 \
flashlinux-0.3.4-RC2.iso-part3 > flashlinux-0.3.4-RC2.iso.
```

Flash Linux has a beautiful Bootsplash and framebuffer theme. It also includes the accelerated NVIDIA driver, which is great if you have a GeForce video card.

Hardware detection also was fantastic. Flash configured all my devices without a hitch.

After you boot and log in, install Flash on the pendrive. You will need two partitions on the pendrive: a boot partition of +4MB and a

second partition of at least 256MB.

The Flash Linux people suggest you set up the partitions with fdisk. Plug in your pendrive and run:

```
fdisk /dev/sda
```

(Again, change sda to the drive designation your computer uses for the pendrive, if it is different from sda.)

Delete every existing partition. Then, add the 4MB partition for the boot partition. Next, create a second partition that uses the rest of the free space on the pendrive. Write the changes and quit fdisk.

Now, download the installation script for USB devices. Download the flash\_key.sh installer file from the Web site, and put it in the root folder of your Flash Linux Live CD.

If the script doesn't see your device, you may need to modify the script. Replace the line:

```
dev=`readlink /${i}|cut -d"/" -f11`
```

in PICKDEVICE with:

```
dev=`readlink /${i}|cut -d"/" -f12`
```

Add execution permissions to the script, and execute it:

```
chmod 755 flash_key.sh
./flash_key.sh
```

Now, follow the easy steps given by the wizard. First, select the correct device to install Flash Linux, then the 4MB boot partition. After that, select the root partition, and install Flash Linux in the pendrive. First, the script erases the pendrive, then mounts it, and finally it copies the apps and data to the pendrive. This last step took more than eight minutes on my machine. Be patient, and after that enjoy Flash Linux.

The highlights of Flash Linux are good speed, thanks in part to the fact that it uses JFFS2 and many applications. Details such as animated cursors and cursor shadows as well as good window decorations, make your Flash Linux Desktop nice.

The only downside to Flash Linux is that it takes many steps to get the pendrive working. Also, I don't know why Flash Linux developers don't include an installer as part of the distribution instead of making you run a script.

## Final Ideas and Impressions

For the desktop user, pendrives and LiveUSB are fantastic. If you have a pendrive, experiment with it—install DSL, SLAX, Flash Linux, Feather, Puppy or other distributions on the hardware.

Your mileage may vary, but I prefer SLAX. The modular nature of SLAX offers a wide range of options and features that a Linux professional should appreciate. SLAX has security modules, the ClamAV antivirus app, a Qt GUI, firewalls and so on. If you work often with security live CDs, pendrives and LiveUSB are also ideal, because with one device, you solve two problems. You can save data, and if you are a developer, have security and development modules on the other side of your pendrive. ■

**Resources for this article:** [www.linuxjournal.com/article/8949](http://www.linuxjournal.com/article/8949).

---

Juan Marcelo Rodriguez has been working with GNU/Linux for many years. He writes articles for magazines, works with a local LUG and also works with LugAR/USLA. He likes to play the keyboard, read, write and listen to music.

# The Ultimate Linux/Windows System

Use cross-platform applications and shared data for the ultimate Linux/Windows system.

KEVIN FARNHAM

I recently converted my Toshiba notebook computer into a dual-boot system, running Windows XP Pro and Ubuntu Linux. I was hoping I'd be able to use cross-platform applications such as Mozilla Firefox, Mozilla Thunderbird, AbiWord, Gnumeric and SciTE transparently, no matter which operating system was currently booted. This article describes the steps I took to make this possible.

## Dual-Boot Computer Configuration for Shared Application Data

In what follows, I assume you already have a dual-boot computer that has a working Linux and Windows operating system installed. You also must have an adequately sized additional disk partition for storing shared application data. This partition must be readable and writable by both operating systems. FAT32 (VFAT) is the logical choice.

My notebook came with Windows XP Pro installed on a 30GB hard drive. The computer was well used, its disk nearly filled, before I decided to convert it to a dual-boot system. I offloaded lots of data, and used the Windows defragment program to reduce my total Windows size below 15GB. Then, I used utilities on the Linux System Rescue CD to resize the original Windows partition and make new partitions as follows:

- Partition 1: Windows NTFS primary partition, 18.5GB.
- Partition 2: Linux ext3 primary partition, 5GB.
- Partition 3: Linux swap partition, 1GB.
- Partition 4: FAT32 partition for shared application data, 5GB.

Making a dual-boot system with only 30GB of total disk space is not ideal. My shared application data partition was 80% full once I loaded my archived e-mail, working documents and various ongoing cross-platform software development projects. For a more ideal setup, I recommend at least 60–80GB of disk space. In that case, I'd allocate 20GB for Windows, 10GB for Linux, 1–2GB for Linux swap and make the remainder the FAT32 shared partition.

## Configuring and Accessing the Shared Disk Partition

Windows views a FAT32 partition as a separate disk drive and assigns it a drive letter. The letter assigned depends on what storage devices are connected to the system—for example, floppy or CD/DVD drives. On my system, Windows identifies the FAT32 partition as drive E:. Use Windows Explorer to verify the Windows drive letter for your FAT32 partition.

When I installed Ubuntu Linux, I selected mounting the FAT32 partition at boot time, using the mountpoint /share. After Linux boots, you can verify that the FAT32 partition is

mounted with the UNIX df command (Listing 1).

Although the /share partition is mounted, there is a problem. By default, the root user owns the /share partition. A standard user will not have read or write permission, and will not be able to run programs that access the shared data. Fortunately, the UNIX mount command provides options for a partition to be mounted with ownership set to a user other than root. This is one method for enabling you to read and write the shared partition using your normal login.

If only one person uses the computer, or only one user needs access to the shared partition, the best plan is to mount the /share partition at boot time, but with your login provided with ownership and full access rights. To configure this, you need to know your user ID and group ID. The /etc/passwd file stores this information. Here's the entry for my user name (kevin) in my /etc/passwd file:

```
kevin@lyratoshibaubuntu:~$
cat /etc/passwd | grep kevin
kevin:x:1000:1000:kevin,,,:/home/kevin:/bin/bash
```

The user ID is the number after the second colon. The group ID is the number after the third colon. The example shows that user kevin is assigned user ID 1000 and group ID 1000 on my system.

Now, you must edit the /etc/fstab file. This filesystem table identifies the filesystems the booting Linux system can expect to see, and instructs Linux on what actions to take for each filesystem. You need to switch to the root user account to edit the file.

First, make a backup copy of the current working /etc/fstab file, so you can revert to that version if something goes wrong. Next, bring the fstab file into an editor, such as vi, emacs, gedit or scite. Find the line for the /share file system, and change the data in the <options> column to defaults,uid=uuuu,gid=gggg where uuuu and gggg are your user ID and group ID from /etc/passwd.

Your finished /etc/fstab file should look something like Listing 2. If multiple user accounts need to access the shared partition, you need a different strategy. One option is not to mount the /share partition at boot time, but instead make a script that users execute to mount the partition giving themselves ownership and full access.

To disable mounting /share at boot time, edit /etc/fstab and place a

Listing 1.

### UNIX df Command Showing Mounted /share Partition

```
kevin@lyratoshibaubuntu:~$ df -k
Filesystem      1K-blocks      Used Available Use% Mounted on
/dev/hda2        5036316    1748816   3031668   37% tmpfs
/dev/shm         184936         0     184936    0% tmpfs
/dev/hda1        184936     12588    172348    7% /lib/modules/2.6.12-9-386/volatile
/dev/hda1       18427896   9955608   8472288   55% /media/hda1
/dev/hda4        4713876    417898   4295978    9% /share
```

Listing 2.

### /etc/fstab File with Boot-Time Mounting of the Shared Partition, Giving Ownership to a Specific User

```
kevin@lyratoshibaubuntu:~$ cat /etc/fstab
# /etc/fstab: static file system information.
#
# <file system> <mount point> <type> <options> <dump> <pass>
proc /proc proc defaults 0 0
/dev/hda2 / ext3 defaults,errors=remount-ro 0 1
/dev/hda1 /media/hda1 ntfs defaults 0 0
/dev/hda4 /share vfat defaults,uid=1000,gid=1000 0 0
/dev/hda3 none swap sw 0 0
/dev/hdc /media/cdrom0 udf,iso9660 user,noauto 0 0
```

# at the start of the /share filesystem line. This makes the line a comment. Then, find the user ID and group ID in /etc/passwd for each user who requires full access to the /share partition. Finally, place a script file similar to the following into the home directory of each user, inserting each user's user ID and group ID after uid= and gid=, respectively:

```
kevin@lyratoshibaubuntu:~$ cat mountShare.csh
sudo mount -t vfat -o uid=1000,gid=1000 /dev/hda4 /share
```

After logging in to Linux, the user opens a terminal window and executes the command script to mount the FAT32 partition with the needed access settings:

```
bash ./mountShare.csh
```

However the shared partition is mounted, you can verify that you have ownership and full access to the /share directory with a long listing of path /:

```
kevin@lyratoshibaubuntu:~$ ls -l / | grep share
drwxr-xr-x 18 kevin kevin 4096 1969-12-31 19:00 share
```

### Using Your Shared Application Data Space

At this point, you are ready to use applications that run on both Windows and Linux to do work on documents stored in your shared application data space. If I'm working under Windows, I store and access my documents using drive E:. Again, the drive letter for the FAT32 partition may be different on your system. If I'm working under Linux, I store and access the same documents in my /share partition.

Before you start editing documents, make sure you have the same version of each application installed on Windows and Linux. Don't just hope there are no configuration or data file structure changes between two different releases of an application.

### Mozilla Suite

I use Mozilla Firefox as my Web browser and Mozilla Thunderbird as my e-mail client. Before I converted my notebook into a dual-boot system, I had run Firefox for a long time, and it had many bookmarks. I also had multiple years of saved e-mail messages. Naturally, with my new dual-boot system, I wanted to run Firefox using all of my previously saved bookmarks, and I wanted to be able to use Thunderbird transparently in Windows and Linux, having full access to all my archived e-mail.

Is this possible? Thanks to the configuration strategy employed by the Mozilla Suite developers, the answer is "yes"! Both Firefox and Thunderbird organize their configurations via profiles. Each profile is

stored in its own subdirectory, which by default is located beneath the top-level configuration directory for the application. The name and location of each profile is stored in an index file named profiles.ini. This structure gives us the flexibility to store the configuration data for any profile in any disk location accessible to the user—for example, on our FAT32 shared application data partition.

Before you make any configuration changes, make sure Firefox and Thunderbird are not running. Then, create a directory on the shared partition where your Mozilla application configuration data will be stored for access by both Windows and Linux. I chose to make a users directory with a subdirectory named kevin, my user name under both operating systems. This is convenient if I decide later to have multiple users on the system. In that case, I'll make a separate

path for each user's unique configuration data, so that the logged in user accesses and maintains his or her own configuration.

Under Windows, the path to my application configuration directory is E:\users\kevin. Under Linux, the path is /share/users/kevin.

### Mozilla Firefox Shared Configuration

For reference, I performed my Firefox shared configuration using Firefox version 1.5. However, the procedures also should work with 1.0.x versions of Firefox.

My Windows Firefox installation contained all my personal bookmarks and other configuration settings, so I reconfigured that Firefox first. In Windows, find the Application Data directory for your user name beneath the Documents and Settings directory. You should see a directory named Firefox that has a subdirectory named Profiles. The Profiles directory will have at least one subfolder. Here's how it looks on my system:

The profiles.ini file tells Firefox where to find its configuration data. Open profiles.ini in an editor, and you should see something like this:

```
[General]
StartWithLastProfile=0

[Profile0]
Name=default
IsRelative=1
Path=Profiles\z9qffpsf.default
Default=1
```

This profiles.ini file identifies my configuration as having a single profile, with the configuration data located in the folder Profiles/z9qffpsf.default, relative to the directory where profiles.ini is located. Looking at the folder tree in Figure 1, you can see the Profiles/z9qffpsf.default folder, with various subfolders. This is the location of all of my unique Firefox configuration information. This is the data I want to be able to access (read and write), whether I am booted in Windows or Linux. In your own Firefox installation, the \*.default folder will have a different name. You need to substitute the name of your own profile directory as you perform the steps described below.

To make my configuration data available to both operating systems, I made a Firefox\Profiles directory beneath my shared E:\users\kevin directory, then copied the original Firefox Profiles/z9qffpsf.default configuration directory to that path. Figure 2 shows the result.

I renamed the original z9qffpsf.default directory on my C: drive to maintain a backup copy in case of unanticipated disaster.

Next, I edited profiles.ini to point to the new location of the Firefox configuration profile. I set the IsRelative flag to zero and the Path location to the shared partition location where I copied the configuration




Figure 1. Windows Firefox Application Data Folder Location and Structure

folder. When setting the path, make sure you use Windows-style backslashes. Otherwise, Windows Firefox won't recognize the new location. Here's my edited profiles.ini file:

```
[General]
StartWithLastProfile=0

[Profile1]
Name=default
IsRelative=0
Path=E:\users\kevin\Firefox\Profiles\ z9qffpsf.default
```

I saved a copy of this file as profiles\_new.ini, so I could return to it in case something went wrong on my first try.




Figure 2. Moved Firefox Configuration Directory

When you've got all of this completed, launch Firefox. If a window pops up asking if you want to import settings from another browser, something is incorrect in your setup. In this case, Firefox will have overwritten the profiles.ini file and created a new default configuration directory. Check your backup copy of your new profiles.ini file and the directory names on the shared partition, make any necessary corrections, re-save your corrected profiles.ini file, and try launching Firefox again. When you have all the configuration elements correct, Firefox launches as it always did, with all of your bookmarks available.

Now, Linux Firefox must be configured to use the same profile. Boot in to Linux and mount the shared partition using one of the described methods. In Linux, Firefox stores the configuration files beneath a user's home in directory .mozilla. Go into this directory, then into the firefox subdirectory, and execute `ls -l`. You'll see a profiles.ini file, the pluginreg.dat file and a configuration profile subdirectory.

To make the Linux Firefox use the configuration profile that was placed onto the shared partition, edit the profiles.ini file. Set the `IsRelative` flag to zero, and set the `Path` to the correct /share location. Here's my modified Linux profiles.ini file:

```
[General]
StartWithLastProfile=1

[Profile0]
Name=default
IsRelative=0
Path=/share/users/kevin/Firefox/Profiles/ z9qffpsf.default
Default=1
```

Start Firefox. If all is correct, you'll see your standard Firefox session with all the bookmarks you originally stored using Windows Firefox available. If this doesn't happen, check the profiles.ini file again, make

certain the /share partition is mounted correctly, with proper ownership and permissions, and verify the exact path to the shared Firefox configuration directory. Replace profiles.ini with your corrected version, and launch Firefox again.

### Mozilla Thunderbird Shared Configuration

The configuration organization for Thunderbird is similar to that for Firefox. For reference, I made my shared configuration using Thunderbird version 1.0.7.

In Windows, find the Thunderbird directory beneath Application Data in the Documents and Settings directory tree for your user name. You might think the Thunderbird directory would be beneath Mozilla, parallel to the Firefox directory, but this wasn't the case on my system. In the Thunderbird directory, you'll see the familiar profiles.ini file and a Profiles folder, as with Firefox.

To make all of your stored e-mail accessible from both your Windows and Linux installations, the configuration folder must be moved to the shared partition. I made a directory \users\kevin\Thunderbird on my shared partition and copied the Profiles directory from the default Windows location into the new shared directory. In my case, the view from Windows Explorer looks like Figure 3.

I renamed my original configuration directory to have a backup and also to be certain that it would no longer be accessed by my operational Windows Thunderbird.

Next, I changed the profiles.ini file to point to the new Thunderbird application data location. My initial profiles.ini looked like this:

```
[General]
StartWithLastProfile=1

[Profile0]
Name=default
IsRelative=1
Path=Profiles\19ximc3t.default
```

I changed IsRelative to zero and set Path to the new Thunderbird location on the shared partition, switching the path directory separators to standard Windows backslashes. Here's my modified file:

```
[General]
StartWithLastProfile=1

[Profile0]
Name=default
IsRelative=0
Path=e:\users\kevin\Thunderbird\Profiles\ 19ximc3t.default
```

When you've done all of this, start Thunderbird. If everything has been modified correctly, Thunderbird starts normally. If the configuration is not correct, Thunderbird will ask about creating a new profile. In this case, cancel and exit the program, check your new profile.ini file and the location of the Thunderbird files on the shared partition. Correct any problems, then run Thunderbird again.

On Linux, you'll find the Thunderbird profiles.ini file in the directory .mozilla-thunderbird, beneath your home directory. Edit profiles.ini to identify the configuration you set up from Windows on the shared partition as the profile Thunderbird should use. Again, set IsRelative to zero and Path to the shared location. Here's my modified Linux Thunderbird profiles.ini file:

```
[General]
StartWithLastProfile=1

[Profile0]
Name=default
IsRelative=0
Path=/share/users/kevin/Thunderbird/Profiles/ 19ximc3t.default
Default=1
```

Launch Thunderbird, and you should have full access to all your e-mail accounts and all the e-mail messages that were saved originally by Thunderbird running on Windows. If Thunderbird asks about creating a new profile, exit and check your work.

### Conclusion

Having a dual-boot Linux and Windows notebook is convenient. The convenience is extended by sharing application data between both operating systems. Being able to run Mozilla Firefox and Thunderbird transparently from both Linux and Windows further enhances a dual-boot system's versatility.

Although a large number of steps are required to create the shared configuration, the individual steps are not difficult for someone familiar with locating, copying and editing files and directory structures in both the Windows and Linux operating systems. ■

**Resources for this article:** [www.linuxjournal.com/article/8954](http://www.linuxjournal.com/article/8954).




Figure 3. Moved Windows Thunderbird Configuration Directory

Kevin Farnham works primarily on software engineering projects involving document indexing, mathematical modeling and simulation, and scientific data acquisition, analysis and presentation. His company, Lyra Technical Systems, Inc., is located in rural Northeastern Connecticut.

# Converting Video Formats with FFmpeg

FFmpeg allows Linux users to convert video files easily between a variety of different formats.

SURAMYA TOMAR

**Today's affordable** digital video cameras have placed the power of digital recording within most people's reach. Unfortunately, this has been accompanied with a corresponding increase in the variety of file formats and codecs available. Some of these formats are more efficient than others, and some are less encumbered by proprietary licensing restrictions. So, having the ability to convert from one format to another is a great help, as you can decide what format you are comfortable with and use that one instead of being restricted to a specific file format.

FFmpeg is a simple and straightforward application that allows Linux users to convert video files easily between a variety of different formats. In this article, I walk you through installing FFmpeg and provide a few instructive examples to demonstrate the range of applications for which it can be used.

## FFmpeg Installation

FFmpeg is an open-source audio and video converter that supports most industry-standard codecs and can convert from one file format to another quickly and easily. It also lets you capture video and audio from a live source and process it.

The source code for FFmpeg is available for download from the project Web site ([ffmpeg.sourceforge.net/index.php](http://ffmpeg.sourceforge.net/index.php)) and at the time of this writing, the latest version available at the site is 0.4.9-pre1.

Once you download the file, extract it using the following command:

```
tar -zxf ffmpeg-0.4.9-pre1.tar.gz
```

This creates a new directory containing the source code for FFmpeg. To install it with the default configuration options, run `./configure` from within the FFmpeg source directory. Once the configuration script finishes, compile it by issuing `make`. Once the compile finishes without any errors, you can install FFmpeg by running `make install` as root.

On the other hand, if you like to have control over what is installed and prefer customizing software installs, you can pass some command-line parameters to the configure script. To see all the options available for the installer, run the following command:

```
./configure --help
```

This command gives you multiple screens of the various settings that can be modified, and you can choose any options you like. The on-screen display does a decent job of explaining what each option does, so I will not go into a lot of detail on this.

I suggest that you enable the following options, but this is not a requirement—feel free to experiment:

- `--enable-mp3lame`: highly recommended—you won't be able to encode MP3s without this. Needs lame to be installed already.

- `--enable-a52`: enables GPLed A52 support, needed for decoding some VOB files.
- `--enable-gpl`: required for the previous component; otherwise, not needed.

As I didn't have lame installed on my system, I ran the following command to configure FFmpeg:

```
./configure --enable-a52 --enable-gpl
```

Once the configuration is complete, read through the output to make sure no errors were generated. Then, run `make`, and go have a drink or something as this might take a little while. Once the system finishes compiling FFmpeg, run `make install` as root to install FFmpeg, and you are done with the installation.

## Basic Usage

Now that you have successfully installed FFmpeg, you can start experimenting with it. The first thing you have to do is choose a video file with which to experiment. As this is your first time with FFmpeg, making a backup copy of this file is highly recommended. You don't want to be responsible for ruining the only copy of a rare video.

This input file most probably has been encoded using a particular codec, but because FFmpeg supports most of the popular formats, we don't need to worry a lot about that. Formats supported by FFmpeg include MPEG, MPEG-4 (Divx), ASF, AVI, Real Audio/Video and QuickTime. To see a list of all the codecs/formats supported by FFmpeg, run the following command:

```
ffmpeg --formats
```

A detailed list of supported file formats is also available at the FFmpeg site.

FFmpeg supports a large list of command-line parameters that controls various settings in FFmpeg. To get a listing of the various options available, run the following command:

```
ffmpeg --help
```

Don't let the multipage listing scare you from using FFmpeg, the basic usage is actually very simple. To convert a file with the default settings, run the following command:

```
ffmpeg -i InputFile OutputFile
```

The `-i` option tells FFmpeg that the filename immediately after it is the name of the file to be used as input. If this option is omitted, FFmpeg attempts to overwrite that file when it tries to create the output file. FFmpeg uses the extension of the output file to try to determine

the format and codec to use, though this can be overridden using command-line parameters (more on this later).

The default settings create an output file that has radio-quality sound (64kbps bitrate) and very bad video quality (200kbps bitrate). Fortunately, these settings can be changed for each encoding, which allows you to choose the quality of each file depending on the need.

To change the audio bitrate, add `-ab bitrate` to the command used earlier, where *bitrate* is the bitrate you want to use. See [www.mp3-tech.org/tests/gb](http://www.mp3-tech.org/tests/gb) for information on the sound quality the various bitrates represent. I prefer to encode files with a bitrate between 128–192kbps depending on my needs, but you can put in a higher value if you so desire. Keep in mind, however, that the higher the bitrate you use, the larger the output file size will be. Also keep in mind that if your source file is encoded in a low bitrate, increasing the bitrate won't accomplish much other than increasing the output file size.

Now, getting a CD-quality audio track for the video doesn't really make sense if the video looks like it was taken using a five-year-old Webcam having a bad day. Thankfully, this problem also is easily solved by adding another parameter to the command line.

To change the video bitrate, add the `-b bitrate` option to the command line. The bitrate here can be any numeric value you like, and I have seen bitrates all the way up to 23,000 (DVD Rips). Although the quality of video encoded with a 23,000kbps bitrate is amazing, the resulting file size of that encoding is also very amazing (a 90-minute video is about 4GB). In my experience, most videos look pretty decent at bitrates between 1,000–1,400, but this is a personal preference, so play with the numbers until you figure out what works for you.

So, to encode a video with a 128kbps audio bitrate and 1,200kbps video stream, we would issue the following command:

```
ffmpeg -i InputFile.avi -ab 128 -b 1200 OutputFile.mpg
```

If you are creating a video CD or DVD, FFmpeg makes it even easier by letting you specify a target type. Then, it uses the target type to calculate the format options required automatically. To set a target type, add `-target type`; *type* can be `vcd`, `svcd`, `dvd`, `dv`, `pal-vcd` or `ntsc-svcd` on the command line. So, if we were creating a VCD, we would run the following command:

```
ffmpeg -i InputFile.mpg -target vcd vcd_file.mpg
```

FFmpeg also has support for encoding audio files. The command to convert audio files is the same as the command to encode video files. To convert a WAV file to a 128kbps MP3 file, issue the following command:

```
ffmpeg -i Input.wav -ab 128 Output.mp3
```

Now, the biggest selling point of FFmpeg is that you can customize it to a level that you are comfortable with. So, if all you want to do is convert from one codec to another, and you don't really care about the advanced features, you can stop reading here and still be able to encode/decode videos. On the other hand, if you like to have more control over the encoding, keep reading as we cover more of the advanced options available in FFmpeg.

There are far too many options available in FFmpeg for me to go over each of them here, so I cover some of the ones I found most interesting and leave the rest for you to explore.

### Forcing the Use of a Particular Video Codec

There are a times when you will want to encode a video using a particular codec and file format. FFmpeg lets you choose the codec with which you want to encode by adding `-vcodec codec` to the command line, where

*codec* is the name of the codec you want to use. So if we want to encode using the MPEG-4 codec at 1,200kbps video bitrate and 128kbps audio bitrate, the command looks like this:

```
ffmpeg -i InputFile.mpg -ab 128 -b 1200 -vcodec mpeg4 OutputFile.avi
```

### Remove the Audio Stream

Let's say you have recorded a video that has a lot of background noise and undesired commentary, so you decide to remove the audio component of the video completely. To accomplish this, all you have to do is add the `-an` option to the command line, and FFmpeg automatically removes all audio from the output. Keep in mind that using this option negates any other option that affects the audio stream.

So, in our example, to remove the audio component, we would run the following command:

```
ffmpeg -i InputFile.mpg -an -b 1200 OutputFile.avi
```

### Remove the Video Stream

Let's say you downloaded a news video from the Net that you want to listen to on your iPod on the way to work, but in order to do that, you have to remove the video component from the output file. FFmpeg allows you to remove the video component of the file completely by adding the `-vn` option to the command line. Using this option negates any other option that affects the video stream.

So, in our example, to remove the video component and save the audio as a 256kbps MP3 file, we would run the following command:

```
ffmpeg -i InputFile.mpg -vn -ab 256 OutputFile.mp3
```

### Choose among Multiple Audio Streams to Encode the Output File

Many DVDs have multiple language tracks available, and you can choose in which language you want to watch the video. Having multiple audio tracks is cool if you speak multiple languages and want to be able to watch videos in multiple languages. However, if you don't speak multiple languages, the extra audio tracks are useless and are taking up disk space.

FFmpeg lets you choose which streams you want to keep and ignore the rest. The command-line parameter that allows you to map streams is called `-map`. So, if in our test file, stream 0 is the video stream, stream 1 is the Spanish audio stream and stream 2 is the English audio stream, and we want to keep the English audio in the output file, we would issue the following command:

```
ffmpeg -i InputFile.mpg -map 0:0 -map 2:1 -b 1200 OutputFile.avi
```

In my experience, stream 0 in most video files is usually the video stream, and the remaining streams are the audio streams available with the video.

### Conclusion

FFmpeg provides a wide range of options for manipulating and converting video files among a variety of formats. For more information, or to download the latest version of FFmpeg for yourself, please refer to the project Web site. ■

---

Suramya Tomar is a Linux system administrator who also likes to program. Visit [www.suramya.com](http://www.suramya.com) for more information on his background.



be **FREE**




Join us at [member.fsf.org](http://member.fsf.org)



# SUSE Rocks, Fedora Locks

Welcome to the 20th century of usability, Fedora Core 5. Too bad this is the 21st.



Nick Petreley, Editor in Chief

I hope you'll forgive me if I can't make good on the name of this column 100% of the time this month, because I will be including a few raves in this rant. I'll start with SUSE 10.0. I have been working on a SUSE 10.0 review for what seems like forever. The way things have been going, SUSE 11.0 probably will ship by the time I'm done.

If you read my rant from last month, you'll know I've upgraded to a dual-core AMD64 machine. I had to reinstall SUSE 10.0 on this new machine to continue my work on the review. I was tempted to call Novell and ask for the AMD64 version of SUSE 10.0, but I went ahead and installed SUSE from the DVD that came in the box they sent. Much to my surprise and delight, it installed an AMD64 version of SUSE. It was only after it installed the 64-bit version that I examined the box carefully. Sure enough, in fine print (at least it's fine print to my 53-year-old eyes), it says that the box includes multiple versions of SUSE, including the AMD64 version.

It gets better. SUSE includes Sun Java 1.5, and SUSE's included i686 version of Firefox

can actually use it without any special work on my part. I had to kludge a 32-bit version of Firefox on 64-bit Kubuntu myself in order to get Java working as a plugin.

I don't want to spend much more time waxing rhapsodic over SUSE, lest I include too much rave in this rant. But although there are a few things I don't like about SUSE 10.0, I would have no trouble recommending it to virtually anyone. I confess, I have never liked SUSE in the past. This version has changed my opinion entirely.

As of a few days before this writing, the final version of Fedora Core 5 became available. I downloaded and installed the AMD64 version. It won't boot. Why? Because Fedora Core configures the GRUB bootloader to boot the partition labeled /. As it happens, I have about three partitions with that same label. I boot the Kanotix live CD, edit the GRUB configuration and `/etc/fstab` to point to the actual partition, and now Fedora boots.

What is it with the Fedora people that they feel compelled to use disk labels instead of partition device names? I realize that this won't be a problem for the average user who uses only Fedora and no other distribution of Linux. But then again, the average user isn't likely to move Fedora to another partition or change the order of drives either, so the advantages of using disk labels will be lost on them. If the Fedora folks are bent on using disk labels, they could at least label the root partition something unique and identifying, like FC564ROOT for the AMD64 version and FC532ROOT for the i386 version.

I'll save you the rant on how disappointed I was with the 64-bit version as a desktop and tell you I decided to go with Fedora Core 5 i386 instead.

This time I edit the GRUB and `/etc/fstab` files after the installation finishes so it will boot the first time. As has been the case for years of using Red Hat distributions, it gets hung up on starting Sendmail about half the

time. Surely I can't be the only person who has experienced this? Why does this problem still exist? Why does Fedora even install Sendmail? There have been superior alternatives for years.

Kudos to Fedora for finally including a software package manager that lets you install packages other than the ones they want you to use. It would be nicer if it were even remotely intuitive, but then the Fedora folks are GNOME-lovers, so making it intuitive would violate the GNOME specification. It would be even nicer if the software updater didn't run so slow that I always assume it is simply frozen and kill it. I always fall back to `yum update`, but even the command-line version of Yum doesn't run, it crawls. Slowly.

The good news about Fedora Core 5 is that it not only includes SELinux (NSA Security Enhanced Linux), but it is preconfigured and enabled by default. It also makes it easy to configure SELinux policies. I love this. If SELinux is important to you, Fedora may be your bag.

But here's my big beef aside from the bonehead disk label problem and other nuisances. Why are the Fedora folks so anal about licenses? There's no Sun Java in Fedora. There's no Flash plugin. I can almost excuse the maintainers for leaving these things out, but Fedora makes no effort at all to make it a no-brainer to add them. How many of you out there really don't want your browser to be able to support Java or Flash? Aren't you going to add these things anyway? So why not make it easy? Ubuntu/Kubuntu makes it easy, and these distros are based on Debian, the most license-anal distro on the planet.

Bottom line—if you're already a Fedora fan, you'll want Core 5. If you use anything else, now's not the time to switch. ■

---

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 23 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.



▲ 72 Port FasTree™ Configuration

Working with PathScale's InfiniPath HTX Adapters, the number of hops required to move MPI messages between nodes is reduced, improving latency. The modular design makes them useful for SDR, DDR and future QDR InfiniBand fabrics, greatly extending their useful life. Please send email to [fastree@microway.com](mailto:fastree@microway.com) to request our white paper entitled *Low Latency Modular Switches for InfiniBand*.

Microway's **QuadPuter®** includes four AMD single or dual core Opteron™ processors, 1350 Watt redundant power supply, and up to 5 redundant, hot swap hard drives—all in 4U. 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™ is also available with an 8-way motherboard and up to 128 GB of memory for power- and memory-hungry SMP applications.



Call us first at 508-746-7341 for quotes and benchmarking services. Find technical information, testimonials, and newsletter at [microway.com](http://microway.com).

◀ QuadPuter® Navion™ with hot swap, redundant power and hard drives and dual core Opterons, offering the perfect balance between performance and density



508.746.7341 [microway.com](http://microway.com)