

LINUX JOURNAL™

Since 1994: The Original Magazine of the Linux Community

APRIL 2006 | www.linuxjournal.com

**Extend Web
Services with
Web Services**

PODCAST AND REEL IN THE BLOGS AND WIKIS

**DAVE TAYLOR
ON WIKIS**

[Edit: Who Cares
What He Thinks?]

**PODCASTING FOR
POCKET CHANGE**

**YOU ARE
NOW ENTERING
BLOG-CITY**

**AGGREGATE BLOGS
AND PODCASTS**

**+ Linux Says It's
Getting Warmer Outside**

AN
SSC
PUBLICATION

USA \$5.00
CAN \$6.50

0 71486 03102 4

04



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
or call Appro Sales at 800.927.5464, 408.941.8100.



LOST

in a sea of server problems?



Get Rescued Today with the AlterPath™ OnBoard!

Download a **FREE** White Paper on Service Processor Management at www.cyclades.com/spm



The Next Big Innovation in Out-of-Band Management

Remote management of servers and other IT assets is critical in today's increasingly complex data centers. Service processor management (HP iLO, IPMI, Dell DRAC, IBM RSA, Sun ALOM) as part of a comprehensive out-of-band infrastructure (OOBI) solution helps enterprises streamline the management of servers from multiple vendors and reduces operating expenses.

The AlterPath OnBoard is a service processor manager that consolidates IP connections, user access and server health information from multiple servers to provide secure, efficient rack-level management with seamless integration into the OOBI.

www.cyclades.com/lja
1.888.cyclades • sales@cyclades.com



CONTENTS

APRIL 2006
Issue 144



ON THE COVER: PHOTO BY CHASE JARVIS

FEATURES

56 WHY I DON'T LIKE [EDIT: GET] WIKIS

Are wikis headed for the Internet graveyard? **Dave Taylor**

58 USING WIKIS AND BLOGS TO EASE ADMINISTRATION

You don't have to use wikis and blogs as they were intended. **Ti Leggett**

62 INSTALLING AND CUSTOMIZING MEDIAWIKI

How to get MediaWiki to strike a happy medium. **Reuven M. Lerner**

66 PODCASTING FOR THE PENGUIN!

Why pay more to create podcasts? We make it simple. **Jon Watson**

70 CREATING A PLANET ME BLOG AGGREGATOR

Reel in those blogs in Planet Me style. **Ben Martin**

76 MOVING TO BLOG-CITY

Want to whip up a fancy blog without getting your hands dirty?
Dee-Ann LeBlanc

82 XOOOPS, YOU CAN DO IT AGAIN AND AGAIN

Blog it yourself with object-oriented extensible content management.
Juan Marcelo Rodriguez

ON THE COVER

- **Extend Web Services with Web Services, p. 22**
- **Dave Taylor on Wikis, p. 56**
- **Podcasting for Pocket Change, p. 66**
- **You Are Now Entering Blog-City, p. 76**
- **Aggregate Blogs and Podcasts, p. 70**
- **Linux Says It's Getting Warmer Outside, p. 90**

The competition doesn't stand a chance.



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

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

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

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



www.coyotepoint.com



CONTENTS

APRIL 2006

Issue 144

COLUMNS

22 REUVEN M. LERNER'S AT THE FORGE

Extending Web Services Using
Other Web Services

32 MARCEL GAGNÉ'S COOKING WITH LINUX

Podcast on the Menu

38 DAVE TAYLOR'S WORK THE SHELL

Calling All Functions, and
Some Math Too!

42 MICK BAUER'S PARANOID PENGUIN

Security Features in SUSE 10.0

50 DOC SEARLS' LINUX FOR SUITS

The Producer Electronics
Revolution

96 NICHOLAS PETRELEY'S ETC/RANT

The Spirit of Open Source

INDEPTH

90 REMOTE TEMPERATURE MONITORING WITH LINUX

How to use Linux to track
global warming (or maybe
just local warming).

Steven M. Lapinskas

IN EVERY ISSUE

8 LETTERS

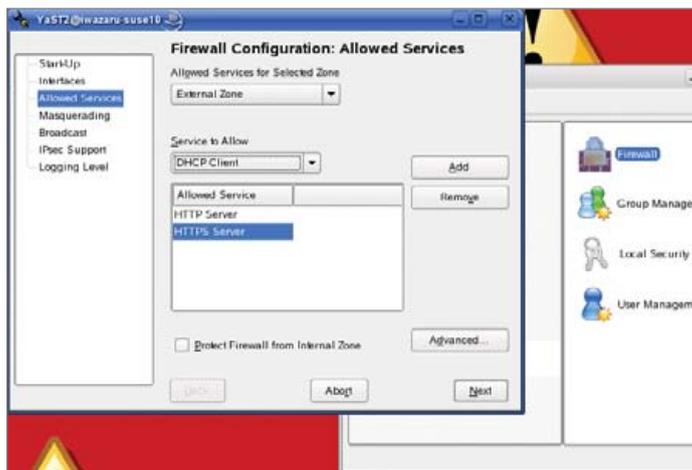
14 UPFRONT

54 NEW PRODUCTS

81 ADVERTISER INDEX



32 This month, Marcel Gagné has various podcasters on the menu, including CastPodder



42 See Mick Bauer's column for more on SUSE's security features and YaST modules.

Next Month

VIRTUAL SYSTEMS

Linux Journal takes a look at the reality behind virtual systems. Want to run a Linux on Linux? User Mode Linux will get you running Debian on Fedora in a safety net. Xen lets you run multiple open-source operating systems simultaneously on the same hardware. Let's not forget about QEMU, yet another virtual system emulation environment. We also look at VMware, the professional virtual system. The best music software still runs on Windows, but we'll tell you how to get it working on Linux too. We'll even tell you how to turn your Linux workstation in to a Sony PlayStation.

Here's something you can't afford to miss. We're bringing aboard John "maddog" Hall with a new monthly column called Beachhead.

There's more. We look at Voice over IP on Linux, how to use Java on embedded systems with GCJ, the GNU Java compiler, and how to manipulate router environments using RTNETLINK.

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™ 8xx series
- 4-P model (B4881-4P):
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

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

LINUX JOURNAL

April 2006, Issue 144

Editor in Chief

Nick Petreley, ljeditor@ssc.com

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

Contributing Editors

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

Proofreader Geri Gale

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

Marketing Manager Rebecca Cassity
rebecca@ssc.com

International Market Analyst James Gray
jgray@ssc.com

Regional Advertising Sales

NORTHERN USA: Joseph Krack, +1 866-423-7722 (toll-free)
EASTERN USA: Martin Seto, +1 905-947-8846
SOUTHERN USA: Laura Whiteman, +1 206-782-7733 x119
INTERNATIONAL: Annie Tiemann, +1 866-965-6646 (toll-free)

Advertising Inquiries ads@ssc.com

Publisher Phil Hughes
phil@ssc.com

Accountant Candy Beauchamp
acct@ssc.com

Linux Journal is published by, and is a registered trade name of, SSC Media Corp.
PO Box 3587, Houston, TX 77253-3587 USA • linux@ssc.com

Editorial Advisory Board

Daniel Frye, Director, IBM Linux Technology Center
Jon "maddog" Hall, President, Linux International
Lawrence Lessig, Professor of Law, Stanford University
Ransom Love, Director of Strategic Relationships, Family and Church History Department,
Church of Jesus Christ of Latter-day Saints
Sam Ockman, CEO, Penguin Computing
Bruce Perens
Bdale Garbee, Linux CTO, HP
Danese Cooper, Open Source Diva, Intel Corporation

Subscriptions

E-MAIL: subs@ssc.com
URL: www.linuxjournal.com
PHONE: +1 206-297-7514
FAX: +1 206-297-7515
TOLL-FREE: 1-888-66-LINUX
MAIL: PO Box 3587, 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., 2825 NW Market Street #208, Seattle, WA 98107. Periodicals postage paid at Seattle, Washington and at additional mailing offices. Cover price is \$5 US. Subscription rate is \$25/year in the United States, \$32 in Canada and Mexico, \$62 elsewhere. POSTMASTER: Please send address changes to *Linux Journal*, PO Box 3587, Houston, TX 77253-3587. Subscriptions start with the next issue. Back issues, if available, may be ordered from the Linux Journal Store: store.linuxjournal.com.

LINUX is a registered trademark of Linus Torvalds.

User Management

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

Wireless AP and Backbone

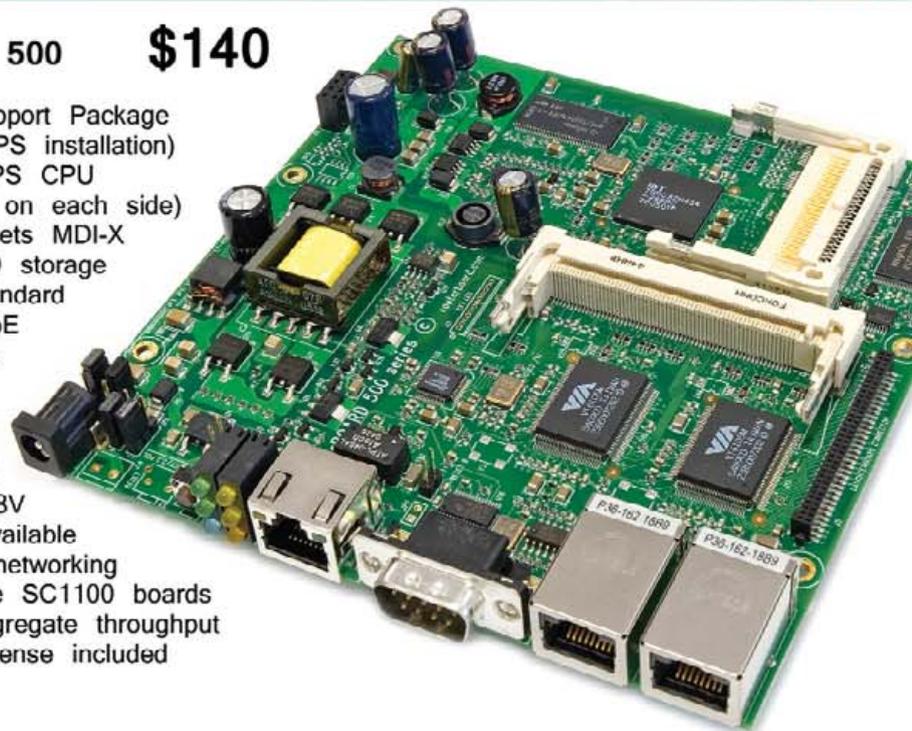
- Wireless monitoring
 - Frequency scanning with detailed report
 - Raw wireless packet sniffer
 - streaming option to Ethereal analyzer
 - option to save to a file format supported by Ethereal
 - Snooper packet inspection
 - analyzes all raw frames received for wireless parameters
 - monitors a single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

- Full 802.11a/b/g support

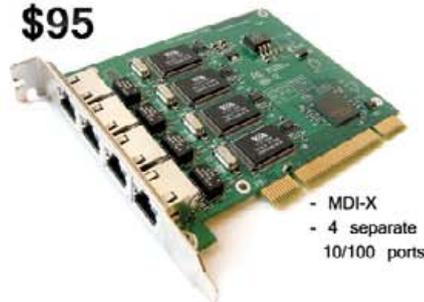
The above is a brief description of a few features, for more information and a fully featured 24 hour demo go to:

RouterBOARD 500 \$140

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



\$95



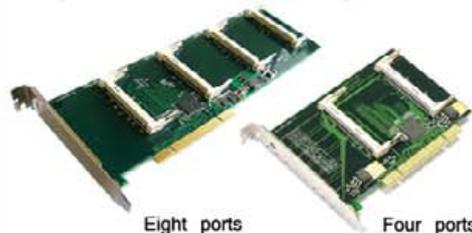
- MDI-X
- 4 separate 10/100 ports

RouterBOARD 44

For the Router Builder !

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

\$120



Eight ports

Four ports

RouterBOARD 11/14/18

Multi radio tower !

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

\$195



Linux Journal Special
Free 64MB SDRAM

RouterBOARD 230

No feature left behind !

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

\$240

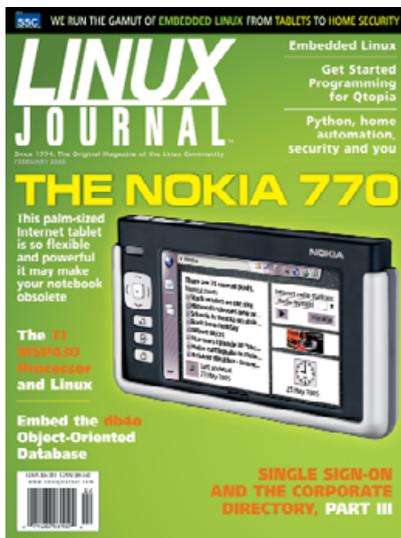
- PoE
- 10-56V input
- 9x 10/100
- 6x mPCI



RouterBOARD 500 & RouterBOARD 564

The Wireless Switchboard !

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



Rant Away

Oh I do love a rant! So, Nicholas, feel free to rant—especially when the rant is as relevant as that in the February 2006 issue of *LJ*. Indeed, let's stop skimming the scum from Redmond's ponds. We need to stop treating MS products as the de facto "gold standard". Of course, this means giving up the unrealistic attempt to make Linux a drop-in replacement for Windows XP, and instead letting Linux stand on its own merits—merits in terms of functionality, initial cost, TCO (as if anyone had any real idea), community, potential for learning and liberty.

You might be interested in the longer, gentler rant by Dominic Humphries "Linux is not Windows", at linux.oneandoneis2.org/LNW.htm, which I give to my students.

Keep ranting! I'd hate to be the only one doing it.

--

Gordon Findlay

No Ho-Hums

I'm a longtime subscriber to *LJ*. I really look forward to every issue, you all do a great job. So, when I read "Nick Gets His Wish" [February 2006], I thought, "okay another 'ho-hum' editor that isn't going to make waves, especially on the first issue". Oh was I in for a surprise on the last page! You really tell it like it is man! Good job. I like you already.

--

Rick Bronson

Your Fine Rant Needs to Cover Even More

There is no reliable, usable way to develop custom small- to medium-sized business applications to run under Linux. Many small- to medium-sized businesses that I develop for would have gladly switched to Linux, until I tell them that they will have to continue running their custom software under a Windows emulator, because there is simply no other feasible way to do it. Then it's, "Why switch to Linux then?" /end/rant, good luck!

--

Terry

According to Evans Data, the vast majority of developers who use Linux as their primary platform for creating small- to medium-business solutions switched to Linux from Windows. They are creating native Linux applications with Eclipse and/or KDevelop and a wide variety of languages and toolkits. There is no need for Windows emulation.—Ed.

Skim Cream Not Scum—OpenOffice.org 2.01

After reading your article about moles in the Open Source community [February 2006], I tried to install OpenOffice.org 2.01 on Windows 98 SE and discovered that it would not load due to the unpacking and loading changes that have been made in XP. I guess OOo 2.01 has an XP loader. So, if you have been using OpenOffice.org on Win 98 SE and try to go to OOo 2.01, you will find that you need XP to unpack OOo 2.01.

I thought you would like to know that your "mole" theory is not too far off.

--

John Farrow

This Is Gonna Be Fun

This [past] month's table of contents [February 2006], not much: MSP 430? Qtopia? Embedded DB40? Augmenting Moore's Law? Case Studies? Nope. The only things of real interest are the Single Sign-On article and the opportunity to drool over the Nokia 770. Next month's Security issue [March 2006] looks interesting, though.

And then I get to /etc/rant. All right. Finally, someone with the guts to say out loud that this slavish emulation of Windows sucks. That flat file text config files with *tons* of comments beat the living crap outta the registry. That we don't *need* a clone of Dot-Net. That OSS has been so busy making sure that we can do everything that the proprietary systems can do that we haven't done anything truly creative in years.

Petreley's EIC of *LJ*. Oh, yeah. This is gonna be fun.

But, according to the LSB, you should have put it under /var.

--

Jeff Simmons

Test This!

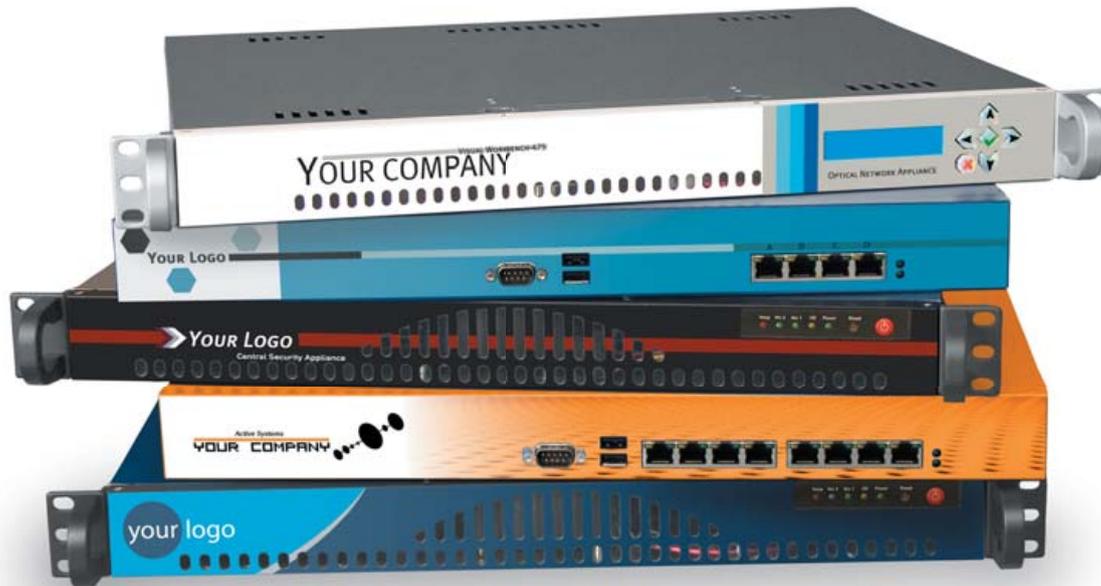
In the January 2006 *LJ* Work The Shell column, Mr Taylor states that "using the [symbol ensures you'll use the built-in version [of test] if available, but explicitly calling test means that you'll likely not have that performance enhancement when running your scripts."

This is totally incorrect. In the AT&T UNIX world, test and [were both built in to the shell starting with System III (circa 1980), and have been ever since. A shell script will never run /bin/test if called as plain test, since built-in commands are always found before external commands. Please pass this clarification on to your readers.

--

Arnold Robbins

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 or call 1-800-681-0016 today.



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

YOUR HIGH PERFORMANCE COMPUTING SOLUTION HAS ARRIVED.

VXRACK™ with the Intel® Xeon™ processor helps you simplify computing operations, accelerate performance and accomplish more in less time.

1 Choose one of the 3 convenient rack sizes

VXR-72

Rack accomodating up to
72 VXBlades/144 Processors
27TB of aggregated Storage
864GB of Global Memory
Power Distribution Included
Patented Architecture
Advanced Cooling System
Integrated InfiniBand Cable Mgnt.



VXR-96

Rack accomodating up to
96 VXBlades/192 Processors
36TB of aggregated Storage
1.15TB of Global Memory
Power Distribution Included
Patented Architecture
Advanced Cooling System
Integrated InfiniBand Cable Mgnt.



For more Information call
or visit us at



VXB-7230NH

Intel SE7230NH Motherboard
800MHz Front Side Bus
Intel® Pentium® D 3.2GHz
1GB DDR2 533 Memory
Single 80GB 7200RPM SATA Drive
One **PCI/Express Slot Available**
Dual 10/100/1000 Intel Lan Port
350W Power Supply

**INTEL® EM64T XEON™
AVAILABLE NOW**

VXB-7520J

Intel SE7520J Motherboard
800MHz Front Side Bus
2 x Intel® EM64T Xeon™ 3.2GHz
2GB DDR2 400 ECC Reg.Mem
Single 40GB 7200RPM ATA Drive
One **PCI/Express Slot Available**
Dual 10/100/1000 Intel Lan Port
500W Power Supply

**2 Choose one or more
type of VXBlade**

**3 Add, Mutiply, That's it.
Easy as 1, 2, 3...**

For example you choose the following: One VXR-96 with
48 Dual Intel® EM64T Xeon™ and 40 Single Intel® Pentium®4.
You take 1 (VXR-96) + 48 (VXB-7520J) + 40 (VXB-7221B)...That's it

VXRACK

THE FUTURE OF CLUSTER TECHNOLOGY

CIARA TECHNOLOGIES...A GLOBAL SOLUTION PROVIDER.

Ciara Technologies is a world-class computer systems manufacturer. Ciara designs, develops, manufactures, markets, services, and supports a variety of computer systems including graphic workstations, rackmount and tower servers, networked storage and the newly acclaimed VXRACK™ Cluster Technology. The company's state of the art supercomputer cluster is based on the Intel IA32 and IA64 architectures and utilizes Linux operating systems. We are proud to be recognized by Intel as an "Intel Premier Provider". Choosing Ciara is choosing a single point of contact for all your IT requirements. All our products are built under the ISO 9001 standards and regulations. The growth of Ciara enabled the company to move its 300+ employees, in February 2003, to an ultra-modern plant of 576,000 sqft.. Ciara now has the capability of producing more than 500,000 units per year.

866-7VX-RACK (866-789-7225)
WWW.VXRACK.COM



Superb Opening Volley

THANK YOU. THANK YOU. THANK YOU for pointing out the absurdity of some the “goings on” in the Linux community of late (re: your /etc/rant on Skim Cream not Scum in the February 2006 LJ).

We love Linux, *Linux Journal* and all things open source. Our CTO and I (both former research Physicists) have watched in amazement as the gatekeepers of things “open source” continue to misplace their angst and focus—with the Sun/Java hangup being front and center.

The *only* people on planet Earth who have not realized that Java *is* open source is a small group of C-exec's at Sun and those Microsoft sycophants in the Linux community you so eloquently point out.

Sure Sun continues to have a hand in the direction of Java—as does Torvalds on the direction of Linux. However, neither Sun nor Torvalds can just do whatever they want—their genies got out of the bottle long ago...

--

Jean

Rants, Maturity and Linux Journal

One of the things I've always liked about *Linux Journal* is the mature, experienced viewpoints. After reading Mr Petreley's rant, I'm wondering if that's going to come to an end very soon. The whole fanatical anti-Microsoft thing is better suited for the “insightful” Slashdot crowd; the real beauty of Linux and open source isn't anti-anything. It's about positive and creative effort, solving problems and providing choice.

--

Eater

Funny you should mention that open source is about choice.—Ed.

Registry Data Is Still Data

I just finished reading your /etc/rant column in the February 2006 issue of *Linux Journal*. Although not generally a Microsoft apologist, there are some things I feel compelled to point out to you. The philosophy behind the Windows registry is not universally wrong. Whether it is a single binary database or a massive XML file is really irrelevant.

--

Tim

XML was a fad choice and unnecessarily obfuscates the information. But the biggest problem I have with the Linux registry is how the data is managed by various daemons.—Ed.

Rant #1

Welcome to *Linux Journal*. I've subscribed for about eight years now and have seen several editors and other staff migrate through. I hope your tenure is fulfilling and mutually beneficial (that is, you and the magazine).

When I read your end-of-issue rant, “Skim Cream not Scum”, I could really relate to it. And, I share your disgust with those that try to emulate Microsoft rather than offer better things.

I use OOo 2.0 only because 1) it converts all my old WP files going back more than a decade, and 2) it can save in the new OpenDocument Format. However, I have also come to like LaTeX (and even came to love LyX for its productivity over writing in Emacs), and that's what I use for articles, reports and my book. I use OOo only to view Word docs sent to me by those still stuck in the Microsoft world (which includes virtually all my clients, agency staff, association copies and relatives) and to send them digital copies of documents they need to read in Word.

Keep ranting about the Emperor not wearing any clothes. Perhaps folks might catch a clue eventually.

--

Rich

Linux at Home—You Conned Me

When I passed the magazine stand at Fry's a few days ago and saw your January 2006 issue, with “Home Projects”, “Linux Phone Home” and especially “At Home with Linux” on the cover, I bought it. I think you conned me. But, hey, it was only five bucks.

What I would like to see is a magazine that tells me how to set up my computer so I can do word processing, data bases and spreadsheets easily. Using Linux. Your magazine does not do that. Not in any part of it.

--

William F. Steagall, Sr.

We have a sister magazine that addresses exactly the sort of thing you describe. It's called TUX (www.tuxmagazine.com). It is available in PDF format.—Ed.

LINUX JOURNAL

At Your Service

MAGAZINE

PRINT SUBSCRIPTIONS: Renewing your subscription, changing your address, paying your invoice, viewing your account details or other subscription inquiries can instantly be done on-line, www.linuxjournal.com/subs. Alternatively, within the U.S. and Canada, you may call us toll-free 1-888-66-LINUX (54689), or internationally +1-713-589-2677. E-mail us at subs@linuxjournal.com or reach us via postal mail, Linux Journal, PO Box 3587, Houston, TX 77253-3587 USA. Please remember to include your complete name and address when contacting us.

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

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

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

ADVERTISING: *Linux Journal* is a great resource for readers and advertisers alike. Request a media kit, view our current editorial calendar and advertising due dates, or learn more about other advertising and marketing opportunities by visiting us on-line, www.linuxjournal.com/advertising. Contact us directly for further information, ads@linuxjournal.com or +1 206-782-7733 ext. 2.

ON-LINE

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

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

SUPERMICR[®]

SAS Server Dual-core Xeon



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-LPZCR1 Zero-Channel RAID support
- ✓ 4 x 1" hot-swap SAS/SATA drive bays w/SES-II
- ✓ 1 slim floppy & 1 slim DVD-ROM drive
- ✓ 600W cold-swap power supply w/I²C
- ✓ 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
- ✓ 6 x 1" hot-swap SAS/SATA drive bays
- ✓ 1 x 5.25" drive bay, 1 floppy drive & 1 slim CD-ROM
- ✓ 500W redundant power supply w/I²C
- ✓ 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-LPZCR1 (Zero-Channel RAID) support
- ✓ 8 x 1" hot-swap SAS/SATA drive bays
- ✓ 90° rotatable module: 2 USB ports, FDD, 2 x 5.25" drive bays
- ✓ 760W triple-redundant power supply w/I²C
- ✓ 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.arrowncp.com

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

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

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

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

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

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

diff -u

WHAT'S NEW
IN KERNEL
DEVELOPMENT

John W. Linville has announced a new **git** repository for his fork of the **Fedora Core** kernels. This repository will take the Fedora kernel releases as its base, and incorporate networking patches from the official kernels on kernel.org. His goals are to allow Fedora users to access bleeding-edge patches they otherwise would have to wait for and to convert the Fedora user community into a testing resource for the kernel (specifically networking) developers.

Eric Dean Moore has assumed maintainership of the previously unlisted **LSI Logic MPT Fusion SCSI Drivers**. These drivers were split recently from a single driver supporting SCSI and Fibre Channel to multiple drivers each supporting a single interface. The irony of this is that although the developers themselves chose to split the driver this way, the MPT Fusion's main claim to fame is that it provides a single interface to both SCSI and Fibre Channel hardware.

Alessandro Di Marco has begun work on **Sluggard**, a filesystem intended to make kernel downloads less painful. As the kernel sources continue to grow, it takes longer to download. Sluggard addresses this by putting rsync at its back end. Users read files on the filesystem with their favorite tools, and Sluggard transparently fetches any changed files just prior to access. In this way, files that aren't needed are never downloaded. This has saved Alessandro more than 200MB of disk space, and he says the potential exists to save even more, depending on what kernel features one needs to compile.

Andrey Volkov has coded up support for the **ST M41T85 real-time clock chip**, basing his work on **Mark A. Greer's ST M41T00 driver**. In fact, there is some talk of merging Andrey's work into Mark's original driver. However, due to many small differences in the hardware, it's not clear what the best approach would be to avoid too many #if statements. It may turn out that merging the two would create more complexity than a unified driver would be worth. But, both Andrey and Mark have indicated that they'd be happy to do a merge if it

seemed like the right way to go.

Since **Sun** released its **ZFS filesystem** source code, **Tarkan Erimer** recently asked if this meant it might be ported to Linux at some point in the future. However, this doesn't seem likely, because Sun used the CDDL license, which allows linking to other code released under nonfree licenses, thus making it incompatible with the GPL. Unless Sun decides to dual-license ZFS under CDDL and the GPL, there is no way a direct port of ZFS to Linux could be legally included in the official tree. It is still possible that someone will do a clean-room reimplementation of ZFS, without looking at any of Sun's sources.

If **Linus Torvalds'** stance on **CVS** was not clear before, it is now. When it was suggested that the kernel include documentation about existing CVS-based kernel repositories, Linus said, "I'd argue against it. CVS is a piece of crap, and anybody who maintains stuff in CVS just makes it harder to ever merge back. That's not just a theory—we've had that situation happen in real life over the years, which is why I definitely don't want to see any external CVS trees given any kind of recognition at all." There you have it.

One of the strengths of the Linux development philosophy is that it won't blindly adhere to standards just because those standards exist. The standards have to make sense as well, and in the case where a good standard has not been followed for some reason, it has to make practical sense to change into conformance with it. This came up recently when **Matthew Wilcox** tried to fix the **NO_IRQ definition** to bring it into conformance with the PCI standard. Linus pointed out that the standard had been ignored by hardware makers for years, and that changing the NO_IRQ definition would break a lot of existing drivers. In fact, the Linux kernel would be clearer and simpler if some change were made to at least bring the various kernel architectures into conformance with each other; and if all drivers could be fixed to reflect this change, it would be the right thing to do. However, Linus rejected that proposal, affirming that the existing design was a mistake, but saying that the fix would break too much outside code.

—Zack Brown

On the Web

Need some advice on figuring out the best way to get the most out of OpenOffice.org? Now that the 2.0 versions are available, more and more people are using this free office suite. You can find your way around OOo well enough if you're moving from another suite, but you'll definitely be more productive with some OOo know-how under your belt. That's where LJ.com columnist Bruce Byfield comes in, with his bi-monthly series OOo Off the Wall. Recent topics include:

- Find and Replace (www.linuxjournal.com/article/8791): how to search for and replace text and strings.
- Master Documents (www.linuxjournal.com/article/8735): how to create and manage large documents comprised of smaller OOo files.
- Domesticating AutoCorrect (www.linuxjournal.com/article/8688): how to set up the AutoCorrect feature so it doesn't drive you crazy—and how to turn it off altogether.

And, for an overview of "What New Users Need to Know about OpenOffice.org" (www.linuxjournal.com/article/8443), Bruce offers advice on on-line help, interface workings, templates and other OOo particulars. In coming months, Bruce will be moving on to Calc, OOo's spreadsheet program.

Regular contributor Colin McGregor is working on a new how-to article for LJ.com readers about the Linux Infrared Remote Control (LIRC) Project (www.linuxjournal.com/article/8811). He'll explain what it is, what hardware/software is needed and how to set up the software. For those of you feeling more adventuresome, Colin also will share some thoughts on building your own infrared hardware.

EmperorLinux

...where Linux & laptops converge

Visit us at the LinuxWorld Expo in Boston from April 4-6, 2006 in booth 1312.



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



The Toucan: 5 lb Linux

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



The Rhino: 7 lb Linux

- Dell Latitude D810/M70
- 15.4" WUXGA screen
- X@1920x1200
- Nvidia Quadro or ATI Radeon
- 1.73-2.26 GHz Pentium-M 7xx
- 40-100 GB HD (7200 rpm)
- 512-2048 MB RAM
- CDRW/DVD or DVD±RW
- 802.11a/b/g wireless, GigE
- Ask about our Sony laptops



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

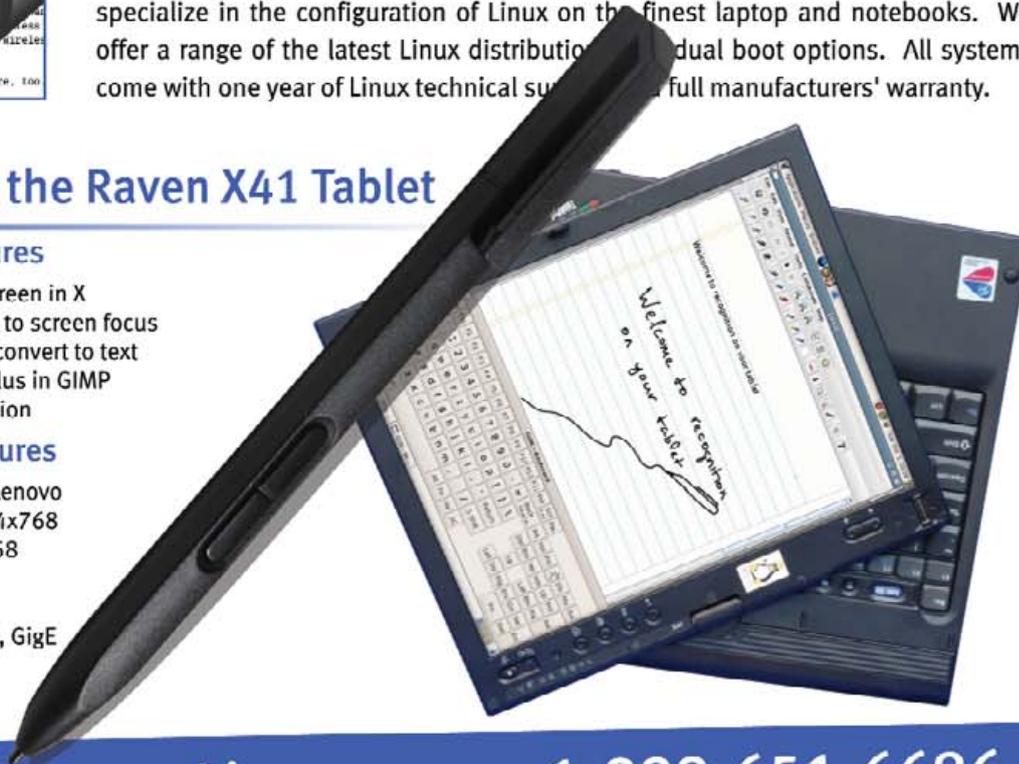
Introducing the Raven X41 Tablet

Linux Tablet Features

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

Linux Laptop Features

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



www.EmperorLinux.com

1-888-651-6686

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

CIARA

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

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

3 Good Reasons to Buy from Ciara 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 sqf. Our systems are build 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, supercomputer clusters as
well as other software and integration servic-
es. All our systems are serviced by Ciara
highly trained and certified technicians and
system engineers. We are an accountable
supplier - One single point of contact for all
your technology needs.

Intel Premier Provider

Ciara have 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.

For all your computer needs visit us at

ciara-tech.com

1-866-789-7225

TECH



\$1099¹

VXPRO-R7230NH **Affordable Starter Server**

Ultra Affordable,
Easy to Service
Base Configuration

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 Ciara

Upgrade Your System

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

\$2955¹

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

Extremely Powerful
Ultra Low Power Consumption
Mid Server Configuration

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 Ciara

Upgrade Your System

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



LJ Index, April 2006

1. Billions of Internet users: **1**
2. Number of computers on the Internet in the US in 1969: **2**
3. Number of years it took for the Internet to reach its first billion users: **36**
4. Estimated year when 2 billion will be reached: **2015**
5. Estimated year when 3 billion will be reached: **2040**
6. Annualized percentage growth rate in Internet use: **18**
7. US percentage of Internet users in 1995: **66**
8. US percentage of Internet users in 2005: **23**
9. Position of South Korea in broadband penetration: **1**
10. Position of China among Internet users under age 30: **1**
11. Millions of Firefox downloads by October 2005: **100**
12. Minimum Firefox market share in November 2005: **10**
13. Number of Netcraft's ten most reliable hosts using Linux: **3**
14. Number of Netcraft's ten most reliable hosts using an open-source OS: **7**
15. Position of Hostway, which runs Linux, among most reliable hosts: **1**
16. Position of Apache among most popular Web servers: **1**
17. Apache percentage of all Web servers: **70.98**
18. Millions of Weblogs as of January 4, 2006: **24.4**
19. Thousands of new Weblogs added each day: **70**
20. Thousands of Weblog posts created per hour: **33**

Sources: 1–7: Jakob Nielsen, Useit.com | 8–12, Mary Meeker, Morgan Stanley | 13–17: Netcraft.com, December 2005 reports | 18–20: Technorati data from January 2006 and November 2005

—Doc Searls

They Said It

Ask the network before you ask the source.

—JEREMIE MILLER, FROM AN IM WITH ME

...consider this a warning shot for anyone who is relying on closed source modules. What you are doing is trying to take from Linux and not give anything back. The GPL explicitly forbids this, and Linux would not be good enough today for you to be using it without that protection. There is a reason why you are wanting to use Linux for your internal use, and why your customers are asking for it.

—GREG KROAH-HARTMAN, REMARKING ABOUT HIS PROPOSED KERNEL PATCH THAT “SIMPLY MARKS ALL PCI FUNCTIONS AS ONLY ABLE TO BE USED BY GPL LICENSED KERNEL CODE” (www.kroah.com/log/2005/11/21#gpl_pci_core)

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

Low Price, Low Power,
High Reliability
for Embedded Systems



TS-7200
shown with
optional A/D converter,
Compact Flash and RS-485

as low as

\$149 qty 1

\$119 qty 100

200 MHz CPU
Power as low as 1/2 Watt

- 3 boards, over 2000 configurations
- PC/104 expansion bus
- Fanless, no heat sink
- up to 256M onboard Flash
- up to 64M SDRAM
- 10/100 Ethernet
- up to 30 DIO ports
- 2 USB ports
- up to 5 COM ports
- Linux, NetBSD
- Real Time extension

Design your solution with one of our engineers

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

See our website for options, peripherals and x86 SBCs



We use our stuff.

visit our TS-7200 powered website at
www.embeddedARM.com
(480) 837-5200

[**UPFRONT**]

USER FRIENDLY by J.D. "Iliad" Frazer



Redraw Your Own Conclusions

Back in July 2001, I had fun looking up Google search results for a bunch of different terms, and listed them in the October 2001 issue of the magazine.

So, we revisited those same items for this issue. Here are the results.—Doc Searls

TERM	JULY 2001	JANUARY 2006
python	2,080,000	91,210,000
active x	2,350,000	107,000,000
gates	3,020,000	78,200,000
kde	3,560,000	29,100,000
gnome	3,720,000	27,100,000
perl	7,650,000	115,000,000
jesus	8,800,000	78,200,000
boy	10,800,000	159,000,000
solution	13,300,000	507,000,000
girl	13,600,000	137,000,000
microsoft	20,200,000	721,000,000
god	24,300,000	172,000,000
sun	25,500,000	523,000,000
sex	28,400,000	211,000,000
linux	31,600,000	420,000,000
business	86,900,000	2,950,000,000
have	231,000,000	3,970,000,000

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



SecureLinx SLC

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

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

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

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

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



LANTRONIX®
Network anything. Network everything.™

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



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





REUVEN M. LERNER

Extending Web Services Using Other Web Services

How to create a useful new Web service by tapping in to the power of two other freely available Web services.

Last month, we looked at the latest incarnation of Web services offered by on-line giant Amazon. Amazon was one of the first companies to embrace Web services, and although some of its newer offerings require payments on a monthly or per-query basis, basic catalog searches are still available free of charge.

If we think of each individual Web service as a function call, we can think of a collection of Web services, such as Amazon Web Services (AWS), as a software library. And although we can certainly create interesting applications with such libraries, it is often useful to create new libraries that sit on top of the existing ones. In many ways, the history of software is the history of creating increasingly powerful abstractions by stacking libraries on top of one another. Outside of the classroom, most of us haven't ever had to implement a sort algorithm or create a buffered I/O library, simply because such things have been written and optimized by previous generations of programmers.

I thus believe that it's useful for us to consider AWS not as a set of routines that we can incorporate into end-user programs, but rather as a set of low-level libraries on top of which we can (and should) create new libraries appropriate for our specific needs.

This month, we look at a simple example of what I mean. The project will reflect my love of books. The Internet has made it difficult for me to stop buying used books, because so many are available at low prices. But, I'm fortunate to be spending several years in Skokie, Illinois, which has an excellent public library. Skokie's library has not only an extensive collection, but it also has a Web-based interface to the book catalog. Our project for this month, thus, is to create a Web service that integrates Amazon's catalog with the information from the Skokie public library. In other words, we're going to write a Web service that itself relies upon another Web service. The input to our service will be an International Standard Book Number (ISBN); the output will be an indication of the book's availability and price at Amazon and the Skokie library.

In some ways, this Web service will duplicate the excellent

Book Burro plugin for the Firefox Web browser, which I often use to find the best bargains. And indeed, Book Burro looks at both bookstores and public libraries in order to find books. I recommend Book Burro to everyone who uses Firefox. But, I believe that building your own simple Web service, even if it duplicates the functionality of another program, is a worthwhile endeavor.

Moreover, Web services have the advantage of being available from any programming language and any application. I can implement my Web service using Ruby, and people will still be able to access it from Java, Python, Perl or virtually any other language. In many ways, this achieves what object broker middleware services like CORBA had promised, only without the baggage that made CORBA a more complex (but arguably more secure and rich) programming platform. It makes a Web service more powerful than a simple software library, because it can be accessed from any platform or language, so long as the requesting computer is connected to the Internet.

Searching the Catalog

In order to integrate an ISBN search for the Skokie library, we're going to need a way to query the library for information about book availability. Unfortunately, my library doesn't have a Web services API for querying its database. But, it does have the next-best thing, namely a simple Web interface that we can query.

There are several ways to look through the output from a Web page. Because many sites now use HTML that can be parsed as if it were XML, we might want to use an XML-parsing library to read through the response from the library's Web site, looking for particular text in specific places.

Much as I might like the idea of such an approach, I'm probably not the only Web developer who takes a more practical, quick-and-dirty look. I have used my library's Web site enough times to know that there is a limited number of responses it might send back to me. As a result, I'll use the reliable, if somewhat stupid, approach of looking for particular cues in the HTTP response.

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

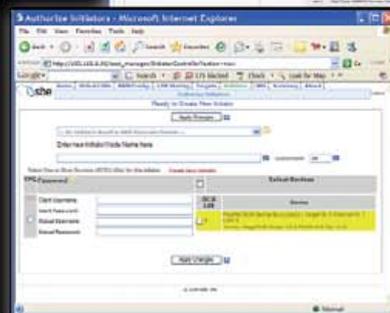
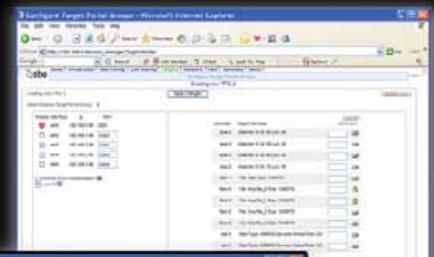


SBE's feature-rich iSCSI software solutions deliver enterprise-level reliability at affordable price points...

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



Our program (skokie-lookup.rb, Listing 1) is written in Ruby, a language I have grown to enjoy more and more over the past few months. We begin by importing the included Net::HTTP module, which defines classes and methods that provide HTTP-based communication.

We then check to make sure that we have at least one

Listing 1.

skokie-lookup.rb

```
#!/usr/bin/ruby

require 'net/http'

if ARGV.length == 0
  puts "#{ARGV[0]}: You must enter at least one argument."
  exit
end

output = ""

# Set up our regular expressions
not_in_collection_re = /class="yourEntryWouldBeHereData"/ix
on_shelf_re = /CHECK SHELF/ix
checked_out_re = /DUE /ix

# Iterate through each of our arguments
ARGV.each do |isbn|

  # Ignore non-ISBN arguments
  if not isbn.match(/[0-9xX]{10}/)
    output << "ISBN #{isbn} is invalid.\n"
    next
  end

  # Ask the library what it knows about our ISBN
  response = Net::HTTP.get_response('catalog.skokie.lib.il.us',
    "/search-54/i?SEARCH=#{isbn}")

  # Check our regular expressions against the HTML response
  if not_in_collection_re.match(response.body)
    output << "ISBN #{isbn} is not in the Skokie collection.\n"
  elsif on_shelf_re.match(response.body)
    output << "ISBN #{isbn} is on the shelf.\n"
  elsif checked_out_re.match(response.body)
    output << "ISBN #{isbn} is currently checked out.\n"
  else
    output << "ISBN #{isbn} response: Unparseable!\n"
  end
end

# Show everyone what we've learned
puts output
```

command-line argument, by looking at the built-in ARGV array. If the length of ARGV is 0, we know we weren't passed any arguments, and we should give the user a brief indication of how the program should be used.

Then, we set up a number of variables that will be needed later on. The output variable is a string to which we will add any output we need to send to the user. We also create three Regexp (regular expression) objects, which we will use in our loop.

Next comes the meat of the program. We iterate over each element of ARGV, first checking that it is a ten-character ISBN containing only numbers and the letter X. We then query the Skokie library's Web site for that ISBN, passing Net::HTTP.get_response the hostname and path to the program we want. The HTTP response, including its headers and body, is then available in our response variable.

Now we compare the response body against our three regular expressions, checking which it matches. Using Ruby's << operator for concatenation, we add an appropriate message to the output variable for each ISBN. Finally, just before the program exits, it gives a full report of ISBNs.

Combining the Search Results

The above program works just fine, and it provides an easier way to query the Skokie library catalog than the standard Web pages. But, I'm interested in knowing how much the book would cost if I were to buy it from Amazon, as well as whether it's available from the library. With all of this information, I can then decide if I want to buy the book, check it out of the library or neither.

Last month, we saw how we could use a REST-style request (that is, HTTP GET with arguments) to retrieve information from Amazon. Now we will write a program that performs that retrieval and then pulls out the relevant XML data.

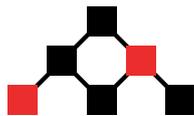
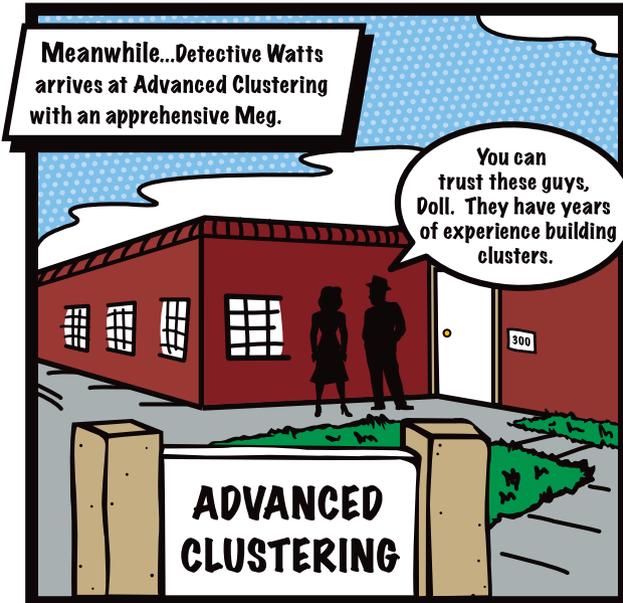
As you might remember, we can retrieve Web services information from Amazon by sending an HTTP request to `webservices.amazon.com`, asking for the document `/onca/xml`, and then specifying the Service, Operation and AWSAccessKeyId name-value pairs. If we are interested in learning about new and used prices for that ISBN, we then pass the ItemId parameter, and indicate that we want the ResponseGroup known as OfferSummary.

Because Amazon returns XML in all of its responses, including those invoked with REST, we can parse through the XML to find the lowest prices for our book. Ruby comes with the REXML-parsing library, which works with XML in a number of different ways; we will use it to scan through Amazon's response for the appropriate code.

Finally, we can rework our existing code, such that it will search the Skokie library for the ISBN and produce a textual summary. Listing 2 contains a program (combined-lookup.rb) that produces such combined output.

combined-lookup.rb begins in almost the same way as skokie-lookup.rb, although it imports the `rexml/document` module along with the `net/http` module. It then iterates through ISBNs that were passed on the command line, ignoring those that don't fit the strict definition.

The Adventures of
DETECTIVE WATTS



advanced clustering
technologies, inc.
Your Turnkey Cluster Provider

www.advancedclustering.com - 1.866.802.8222 - info@advancedclustering.com

The main addition to this program begins with the creation of a string named `amazon_params`. In theory, we could have built this string in a number of different ways, many of them less complicated than the combination of methods I chose.

But, I felt that using a hash in this way would make it easier to modify the code later on, even if it requires a bit more time to understand at first.

The basic idea is as follows: we create a hash, in which

Listing 2.

combined-lookup.rb

```
#!/usr/bin/ruby

require 'net/http'
require 'rexml/document'

if ARGV.length == 0
  puts "#{$0}: You must enter at least one argument."
  exit
end

output = ""

# Set up our regular expressions
not_in_collection_re = /class="yourEntryWouldBeHereData"/ix
on_shelf_re = /CHECK\S+SHELF/ix
checked_out_re = /DUE /ix

# Iterate through each of our arguments
ARGV.each do |isbn|

  # Ignore non-ISBN arguments
  if not isbn.match(/[0-9xX]{10}/)
    output << "ISBN #{isbn} is invalid.\n"
    next
  end

  output << "ISBN: #{isbn}\n"

  # -----
  # Amazon
  # -----

  # Put together an Amazon parameter string
  amazon_params = {'Service' => 'AWSCommerceService',
                  'Operation' => 'ItemLookup',
                  'AWSAccessKeyId' => 'XXX',
                  'ItemId' => isbn,
                  'ResponseGroup' => 'Medium,OfferFull',
                  'MerchantId' => 'All'}.map {|key,value| "#{key}=#(value)"}.join("&")

  # Ask Amazon what it knows about our ISBN
  amazon_response = Net::HTTP.get_response('webservices.amazon.com',
                                           '/onca/xml?' << amazon_params)

  xml = REXML::Document.new(amazon_response.body)

  # Get the lowest new, used, and collectible prices

  new_price =
  xml.root.elements["Items/Item/OfferSummary/LowestNewPrice/FormattedPrice"]
  if new_price.nil?
    output << "\tNew: None available\n"
  else
    output << "\tNew: #{new_price.text}\n"
  end

  used_price =
  xml.root.elements["Items/Item/OfferSummary/LowestUsedPrice/FormattedPrice"]
  if used_price.nil?
    output << "\tUsed: None available\n"
  else
    output << "\tUsed: #{used_price.text}\n"
  end

  collectible_price =
  xml.root.elements["Items/Item/OfferSummary/LowestCollectiblePrice/FormattedPrice"]
  if collectible_price.nil?
    output << "\tCollectible: None available\n"
  else
    output << "\tCollectible: #{collectible_price.text}\n"
  end

  # -----
  # Library
  # -----

  # Ask the library what it knows about our ISBN
  library_response = Net::HTTP.get_response('catalog.skokie.lib.il.us',
                                           "/search~S4/i?SEARCH=#{isbn}")

  # Check our regular expressions against the HTML response
  if not_in_collection_re.match(library_response.body)
    output << "\tLibrary: Not in the Skokie collection.\n"
  elsif checked_out_re.match(library_response.body)
    output << "\tLibrary: Checked out.\n"
  elsif on_shelf_re.match(library_response.body)
    output << "\tLibrary: On the shelf.\n"
  else
    output << "\tLibrary: Unparseable response\n"
  end

  # Show everyone what we've learned
  puts output
end
```

Listing 3.

xmlrpc-lookup.rb

```
#!/usr/bin/ruby

require 'net/http'
require 'rexml/document'
require 'xmlrpc/server'

# Set our regular expressions
not_in_collection_re = /class="yourEntryWouldBeHereData"/ix
on_shelf_re = /CHECK\s+SHELF/ix
checked_out_re = /DUE /ix

# -----
# XML-RPC
# -----

# Start an HTTP server on port 8080, to listen for clients
server = XMLRPC::Server.new(8080)

server.add_handler(name="atf.books",
                  signature=['array', 'array']) do |isbns|

  output = [ ]

  # Iterate through each of our arguments
  isbns.each do |isbn|

    isbn_output = {'ISBN' => isbn}

    # Ignore non-ISBN arguments
    if not isbn.match(/^[0-9xX]{10}$/)
      isbn_output['message'] = "ISBN #{isbn} is invalid."
      output << isbn_output
      next
    end

    # -----
    # Amazon
    # -----

    # Put together an Amazon parameter string
    amazon_params = {'Service' => 'AWSECommerceService',
                    'Operation' => 'ItemLookup',
                    'AWSAccessKeyId' => 'XXX',
                    'ItemId' => isbn,
                    'ResponseGroup' => 'Medium,OfferFull',
                    'MerchantId' => 'All'}.map {|key,value| "#{key}=#{"value"}}.join("&")

    # Ask Amazon what it knows about our ISBN
    amazon_response = Net::HTTP.get_response('webservices.amazon.com',
                                             '/onca/xml?' <<amazon_params)

    xml = REXML::Document.new(amazon_response.body)

    # Get the lowest new, used, and collectible prices
    new_price =
    xml.root.elements["Items/Item/OfferSummary/LowestNewPrice/FormattedPrice"]
    if new_price.nil?
      isbn_output['New'] = "None available"
    else
      isbn_output['New'] = new_price.text
    end

    used_price =
    xml.root.elements["Items/Item/OfferSummary/LowestUsedPrice/FormattedPrice"]
    if used_price.nil?
      isbn_output['Used'] = "None available"
    else
      isbn_output['Used'] = used_price.text
    end

    collectible_price =
    xml.root.elements["Items/Item/OfferSummary/LowestCollectiblePrice/FormattedPrice"]
    if collectible_price.nil?
      isbn_output['Collectible'] = "None available"
    else
      isbn_output['Collectible'] = collectible_price.text
    end

    # -----
    # Library
    # -----

    # Ask the library what it knows about our ISBN
    library_response = Net::HTTP.get_response('catalog.skokie.lib.il.us',
                                              "/search-S4/i?SEARCH=#{isbn}")

    # Check our regular expressions against the HTML response
    if not_in_collection_re.match(library_response.body)
      isbn_output['Library'] = "Library: Not in the Skokie collection."
    elsif checked_out_re.match(library_response.body)
      isbn_output['Library'] = "Checked out."
    elsif on_shelf_re.match(library_response.body)
      isbn_output['Library'] = "On the shelf."
    else
      isbn_output['Library'] = "Unparseable response."
    end

    output << isbn_output
  end

  output
end

server.serve
```

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

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

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



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

the keys are the AWS REST parameter names, and the values are the corresponding parameter values. In order to get these parameters into the standard format of `param1=value1¶m2=value2`, we use `map` to create an array from the keys and values of the hash. Our array will contain strings, each of which is in the `param=value` format, joined together with an equal sign. Finally, we use `join` to combine all of those pairs with `&` signs between them, producing a string that we assign to `amazon_params`.

With our parameters in place, we use `Net::HTTP.get_response`, just as we did before in `skokie-lookup.rb`. The hostname will be different, and the requested URL on that host will also be quite different, incorporating the parameters that we just assigned to `amazon_params`. But, the request is sent in the same way, and we retrieve the response in the same way as well.

However, whereas the Skokie library sends its response in HTML, Amazon replies using XML. So, we fire up `REXML`, creating a new instance of `REXML::Document` with the contents of the Amazon response. We then use the `elements` method on the response's root node to find the lowest new, used and collectible prices. (Amazon provides each of these prices separately, which I admit is a bit annoying.) If the text within that node is nil, no such price exists, and we indicate that to the user. Otherwise, we can assume we got a price back—and a price formatted with a dollar sign and decimal point, at that—and we display it for the user.

Creating a Web Service

Now that we have created a combined lookup tool, how can we turn it into a Web service? (For the purposes of simplicity, I'm going to use XML-RPC. It would be equally valid to use SOAP or even to look for REST parameters.)

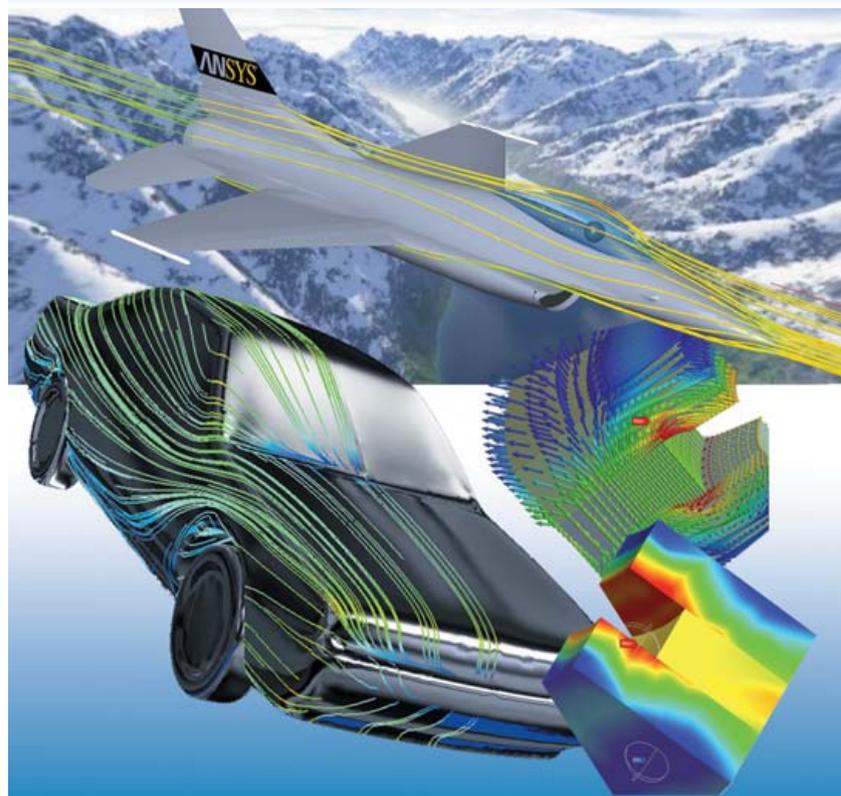
The answer is easier than you might think. We will need to modify the program to take its inputs from the Web instead of ARGV. We also will need to send the output over the XML-RPC, back to the client that sent the original request.

But the end result, as you can see in Listing 3, is not terribly different from what we had in Listing 2. And because it operates as a Web service, we can now incorporate its results into new programs that we might write. Better yet, we can create new Web services that use this service as an underlying foundation, thus stacking the functions even deeper, into even more useful libraries.

Listing 3 begins by creating a new instance of `XMLRPC::Server` on port 8080. It then adds a new handler, which we call `atf.books`, and which both accepts an array as input and returns one as output. Using Ruby's block notation, the handler then iterates over each ISBN that it receives via the XML-RPC method call.

The rest of the program is largely the same as `combined-lookup.rb`, with the exception of the output. Output to an XML-RPC method call, at least in this Ruby library, is accomplished by placing the output in the final line of the block. Because we plan to return an array, we need to create and populate the array. We thus define output variable as an

64-bit ANSYS® Compiled With PGI



For more than 30 years, companies have relied upon ANSYS engineering simulation solutions to deliver innovative, winning products to market better, faster and cheaper. Today, organizations in a wide range of industries turn to ANSYS solutions for 64-bit systems to bring their product development to the next level of innovation.

Visit www.ansys.com/pgi to learn more about ANSYS simulation solutions for 64-bit systems.

ANSYS software for AMD Opteron processor-based systems is built using *PGI Compilers and Tools*.



www.ansys.com

Listing 4.

xmlrpc-lookup-client.rb

```
#!/usr/bin/ruby

require 'xmlrpc/client'

# Get the ISBNs from the command line
isbn = ARGV

# Connect to the server
server = XMLRPC::Client.new2("http://127.0.0.1:8080/", nil, 120)

# Send the ISBNs, and catch any faults that we find
begin
  results = server.call("atf.books", isbn)
rescue XMLRPC::FaultException => e
  puts "Error:"

  puts e.faultCode
  puts e.faultString
end

# Display the results!
results.each do |result|
  result.each do |key, value|
    if key == "ISBN"
      puts "ISBN: #{value}\n"
    else
      puts "\t#{key}: #{value}\n"
    end
  end
end
```

empty array and add one element to it for each ISBN we check. Each element of that array then will be a hash (known as a struct in XML-RPC jargon), with the ISBN key pointing to the book's ISBN, and the New, Used and Collectible keys pointing to the prices retrieved from Amazon.

The server program then concludes with a call to `server.serve`, starting an infinite listener loop for a simple HTTP server.

To test this program, you need an RPC client; a simple one is shown in Listing 4 and takes its arguments from the command line. You'll notice that we use Ruby's exception-handling mechanism to watch for potential problems. If there is an error on the server, we can trap it and print a useful debugging message.

Conclusion

Seasoned programmers rarely implement everything themselves. The days in which every application needed its own video and printer drivers, to say nothing of a filesystem or operating system, are long behind us. Instead, we now have hierarchies of software libraries, with each library making use of lower-level data and functions and also performing similar tasks for higher-level libraries.

Web services haven't changed the need for building new libraries on top of old ones. Indeed, we can expect to see an explosion of such new libraries in the future. The difference is that new libraries will often be based on Web services, which provide platform and language independence. We will see basic, middleware and high-level Web services, available from anywhere on the Internet and callable from any operating system or language. This month, we looked at one way in which we can create a new Web service out of an old one. Each call to our `xmlrpc-lookup` server fired off a query to Amazon's Web services. Information from Amazon was then combined with another data set, with results that are useful to anyone living in Skokie, Illinois. We can expect to see similar aggregating Web services in the future, both free of charge and for pay.

Resources for this article: www.linuxjournal.com/article/8828. ■

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.

WARNING:

NIGHTSTAR LX™ IS HABIT-FORMING, AND MAY CAUSE EXTREME FEELINGS OF EUPHORIA.



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



Download a **FREE TRIAL** at
www.ccur.com/nightstar/LX2
 800-666-4544 • 954-974-1700
 Email us at nightstar@ccur.com

© 2006 Concurrent Computer Corporation



Big business power, small business price tag

Robust ServersDirect Systems powered by the 64-bit Intel® Xeon™ Processor meet your needs today and grow with you into the future

2U Dual Core + SAS Servers



SDR-2115A 2U SAS Advanced Application Server

Provides ideal solution for running multiple applications or functions simultaneously while designed for high performance enterprise requirements.

- 2U rack designed with 6 x 1" Hot-swap SAS / SATA-II Drive Bays
- Dual Xeon™ Dual Core 800MHz processors supported
- Supermicro® board with Intel® E7520 Chipset
- Support up to 16GB DDRII 400 SDRAM memory
- Maximum drive capacity support up to 882GB (SAS) or 3TB SATA-2
- ATI Rage XL SVGA PCI video controller with 8 MB of video memory
- Adaptec AIC9410 8-Port Serial Attached SCSI (SAS) Controller
- RAID 0, and 1 support
- Dual Intel® 82541PI Single-Port Gigabit Ethernet Controllers

1U Appliance Server SDR-1110T

WORLD'S ONLY Sub-Thousand!
Dual Core Server in 1U Solution



- Intel® Pentium® DUAL CORE 2.8 2X1M 800MHZ Processor
- 1U Chassis with 300W power supply
- Supermicro® server board with Intel® E7230 (Mukilleo) Chipset
- Kingston 512MB 667MHz DDR2 ECC RAM (2x256MB)
- 1pc x Seagate 80GB SATA 7200RPM hard drive
- 2 x 1" Hot-swap SATA drive bays
- Integrated ATI Rage XL SVGA PCI video controller
- 2x Intel® 82573V (Tekoa) PCI-e Gigabit LAN Ports
- Built-in SATA ICH7R Controller 4x SATA (3 Gbps) Drive with RAID 0, 1, 5, 10 Support

\$999

3U Xeon® Storage Server SDR-3309T

Highly manageable storage server
supporting up to 8TB of SATA hot-swap
Storage



- Intel® Xeon™ Processor 2.8Ghz with 800FSB 1MB Cache (Dual Processor Option)
- Intel® Extended Memory 64 Technology
- 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)
- Intel® Extended Memory 64 Technology
- 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)
- Intel® Extended Memory 64 Technology
- 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.





MARCEL GAGNÉ

Cooking with Linux Podcast on the Menu

If you don't get a buzz off the wine, you will surely get high on these podcast download tools.

Relax, François! I wasn't trying to scare you. I pulled your earphones out because it is nearly time for our guests to arrive, and I couldn't seem to get your attention any other way. What were you listening to anyhow? The *New Scientist* podcast? Very nice, *mon ami*. I did not realize that you were fascinated by science stories. I must admit, that particular podcast is one of my favorites as well. I regularly download several programs including that one. Sometimes I listen to them on the Linux system where I have downloaded them. Other times, I copy them to my handheld and take them on the road. Podcasts are a great diversion on trains or while sitting in waiting rooms. Best of all, there are tons of great programs out there, from pros and amateurs alike with topics covering pretty much anything you can think of.

Quoi? It takes a while, you say? You don't really manually download all your podcasts, do you? Right-click and Save As off the Web page? Really, François, there are much better ways to do this. Luckily, those helpful programs are the subject of tonight's menu. But for now, I see that our guests have arrived and we must attend to them *immédiatement!* Welcome, everyone, to *Chez Marcel*, home of fine wine, delectable Linux fare and, of course, the most excellent clientele. Please, allow me to help you to your tables. Make yourselves comfortable. François, to the wine cellar! Head to the north wing and bring back the 2002 New Zealand Pinot Noir, the Marlborough, please.

François and I were just discussing podcasts and the need for some way to download and collect them effectively. Chasing down podcasts, looking for new programs and downloading them can take a frightful amount of time. That's why podcast aggregators are such a great idea and why so many clients exist. Some graphical clients are rich and complex, like the wine—we'll look at one shortly. Others are much simpler. One of the simplest podcast aggregators I've seen, Linc Fessenden's BashPodder, is nothing more than a bash shell script and a tiny one at that. Head over to the BashPodder Web site (see the on-line Resources), and download the `bashpodder.shell` script and the `parse_enclosure.xml` file. When you get the two, copy them into the folder of your choice. Then, have a look at the script. At less than 1K, you won't believe how simple it is.

Excellent, François. That was very fast. Please pour for our guests. Enjoy, *mes amis*.

A sample configuration file, `bp.conf`, is also available for

download from the BashPodder site. The sample file includes some of Linc's favorite podcast feeds, so feel free to download it if you are curious as to his tastes, but the format is simple enough that you won't need it. `bp.conf` is simply a collection of feed links. For instance, here's a file with three feeds, one for the *New Scientist* podcast, Brian Ibbott's *Coverville* and the BBC's excellent *In Our Time*.

<http://www.newscientist.com/podcastfeed.ns>

<http://www.coverville.com/index.xml>

<http://www.bbc.co.uk/radio4/history/inourtime/mp3/podcast.xml>

I store this file in the same directory from which I run the `bashpodder.shell` script file:

```
sh bashpodder.shell
```

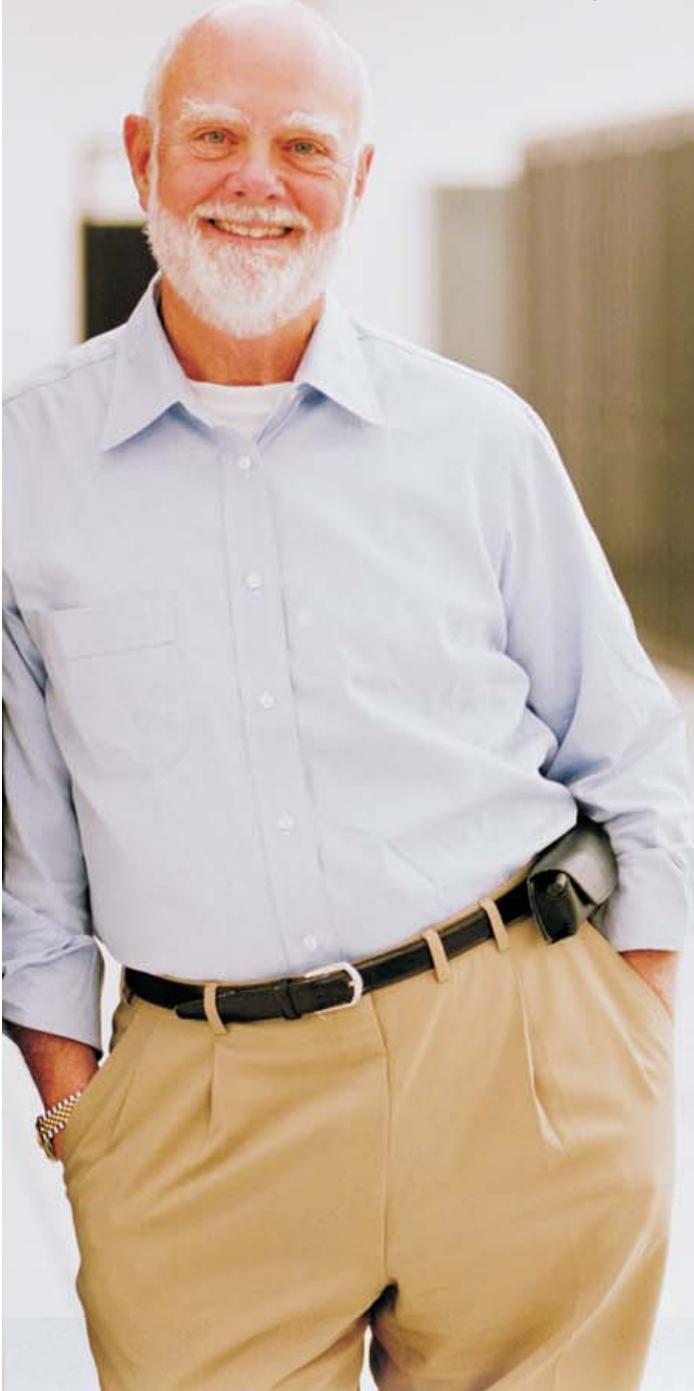
The script creates a folder with the date of the run as its name (for example, 2006-01-04). BashPodder runs silently and downloads all of the current podcasts to that folder. Set the program to run in a cron job, and you will always be right up to date. Be warned, however, that BashPodder will download all of the current downloads for the particular feed. There's no picking and choosing.

Eskild Hustvedt's GoldenPod is also a command-line driven program that is simple to use and easy to set up in a cron job for unattended downloads. This one uses Perl instead of bash and adds some useful features that aren't present in the much simpler BashPodder. GoldenPod can run silently (with the `-s` command switch), but it also can report on its progress as it works. You can elect to download only the most recent podcast in a series rather than the entire collection (which can sometimes be huge). GoldenPod also can list available podcasts without downloading, copy files to a portable player, clean up old podcasts and provide statistics on how much space your podcasts are taking up. All this for less than 32K.

Get your copy of GoldenPod from the Web site (see Resources) and extract the tarball into the directory of your choice. If you simply run `./goldenpod` from the folder in which you extracted it, the program will report that your podcast list is empty, but it also will create two new directories for you in your home folder. The first is called `.goldenpod`, which contains your configuration files and podcast list. The second is called `Podcasts`, and this is where your podcasts will be down-

Big business power, small business price tag.

Run your business on the Processor many of the world's largest companies choose. Robust ZT X9000 series systems powered by the Intel® Xeon™ Processor meet your needs today and grow with



ZT 2U Server X9555

Dual Intel® Xeon® Processors 3 GHz
(2MB L2 Cache, 3 GHz, 800MHz FSB)

- Intel® E7520 Chipset Server Board
- 1GB DDR2 400MHz ECC/ Reg SDRAM (Up to 16GB)
- 2 x 36GB 10,000rpm SATA Hard Drive (RAID 1 for O/S)
- 6 x 150GB 10,000rpm SATA Hard Drive (RAID 5 for DATA)
(Convertible to Seagate® 15,000rpm SAS Hard Drive)
- 8 x 1" Hot-swap SATA Drive Bays
- 64bit High Performance SATA2 / SAS RAID Controller
- 2 x Intel® Gigabit Ethernet
- 3U Rackmount Chassis w/ 760W Redundant Power Supply
- SuperDoctor III Server Management Software
- 3-Year Limited Warranty



\$4,999

Find out how ZT insider Program can help maximize your business solution

Go to

ztgroup.com/go/linuxjournal

Call

866-ZTGROUP (984-7687) Promote Code # IJ0406



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

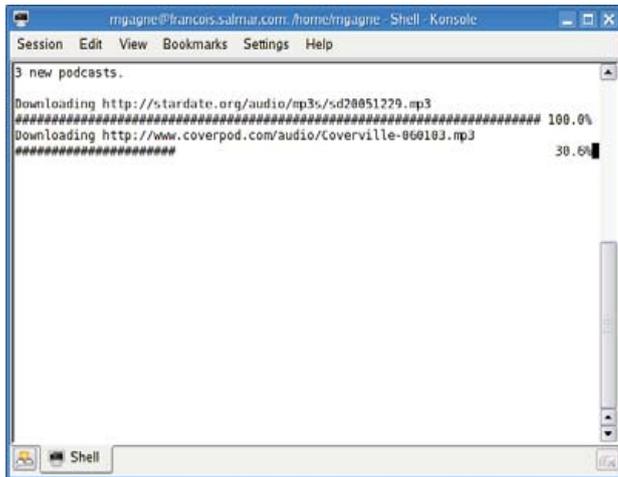


Figure 1. GoldenPod is a simple, command-line-driven podcast aggregator.

loaded. In the `~/goldenpod` folder, there are two files: `goldenpod.conf` and `podcasts.conf`. The format of the `podcasts.conf` file—a text file that you can edit using your favorite editor—is the same as BashPodder's `bp.conf` file, essentially a list of podcast feed addresses.

Once you have added feeds, simply run the program again to start the download. You may want to try `goldenpod --dry-run` to see what you are in for or `goldenpod -f` to download only the latest. You then can see the progress of your downloads on your terminal (Figure 1).

When the download is done, you'll find `m3u` playlist files in each of the podcast directories corresponding to your feeds (for example, `~/Podcasts/catalog/showname/showname.m3u`). Sure, it's all command line, but once again, this is a great program to run in a cron job, quietly in the background, perhaps while you sleep. From time to time, it is also a good idea to run the program with the `--stats` option to see how many programs you have and how much space they are taking up.

Of course, there are graphical podcast aggregators for Linux as well, which, by their nature, expand upon the feature sets of the two I've mentioned. The perfect podcast aggregator would, for starters, contain its own directory of podcasts, categorized into an easy-to-search list of topics. Complement that with an easy one-button system to add feeds, a new show scanning feature so you don't have to go looking for the latest shows and a system of automatic scheduled downloads running in the background, and you've got the makings of a great program. As it turns out, Scott Grayban's CastPodder does all those things. Your first step toward podcast Nirvana is to head over to the CastPodder Web site (see Resources) and download a copy of the latest program. There, you'll find packages for Mandriva and Debian, as well as tarred and gzipped source packages.

When you start CastPodder for the first time (command name: `CastPodder`), it creates a folder and subfolder called `iPodderData/downloads` in your home directory. You can override this default by selecting Preferences from the File menu. The CastPodder interface itself has the usual menu bar running along the top offering access to all the program's features. There are four main tabs labeled Downloads, Subscriptions, Podcast Directory and Cleanup. If it isn't already selected when you start, click the Podcast Directory tab, and you'll see a list of folders describing various directories and top podcast collections available. Before you go hunting for interesting shows, click the Refresh icon on the far left of the icon bar. That's the best

way to make sure you are working from a recent list.

Now, let's get back to those podcast directories. Perhaps the most interesting here is the iPodder.org : Podcasting Central folder, because it contains the massive list of podcasts from iPodder.org. Click the small arrow next to the folder to list the categories, subcategories and, finally, all the various podcasts that correspond to that category. To subscribe to a podcast, click the title and then click the Add button. If the podcast you want isn't listed here, you can add it manually by clicking Tools on the menu bar and selecting Add a Feed. As you subscribe to various podcasts, subdirectories with the names of those podcasts will be created in `~/iPodderdata/downloads`. For instance, I have folders called Coverville, New Scientist Podcast and Quirks and Quarks from CBC Radio among others. Each folder holds that particular show's podcasts, so make sure you have a lot of disk space.

Let's go back to the Subscriptions tab (Figure 3). A list of all the podcasts to which you have subscribed is in the top half of the main window. Click any of these shows, and the bottom pane lists the various episodes, the size of the show

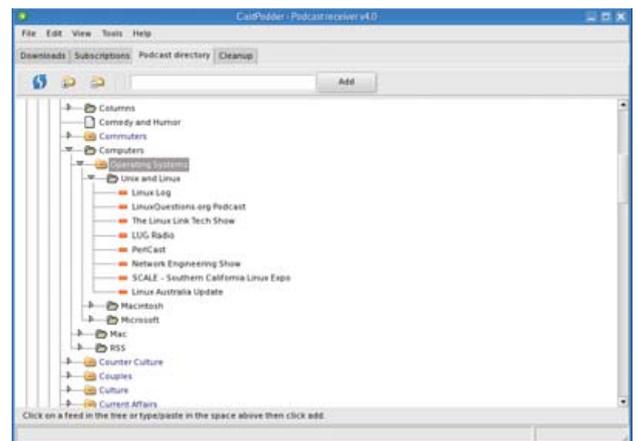


Figure 2. No need to go hunting through broadcast directories—CastPodder provides its own extensive list.

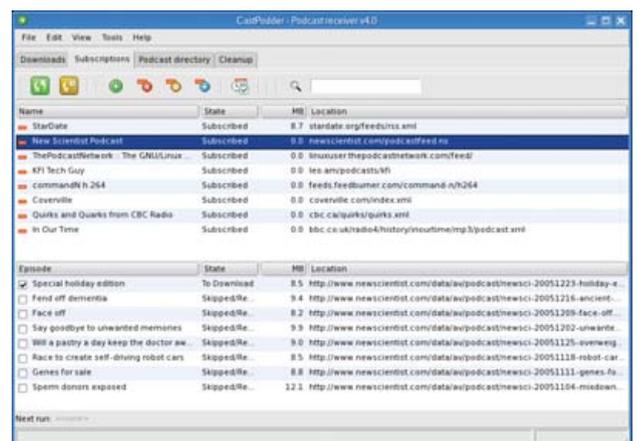


Figure 3. Subscription lists allow you to individually select episodes for download.



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

AoE is your answer!

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

Ethernet Storage – without the TCP/IP overhead!

Unlimited expandability, at the lowest possible price point!!

You want more storage..... you just buy more disks – it's that simple!!!



EtherDrive® SATA+RAID 3U

RAID enabled 3U appliance with 15 slots for hot swap SATA disks

Check out our other Storage Appliances and NAS Gateway



www.coraid.com
info@coraid.com
1.706.548.7200



The Linux Storage People

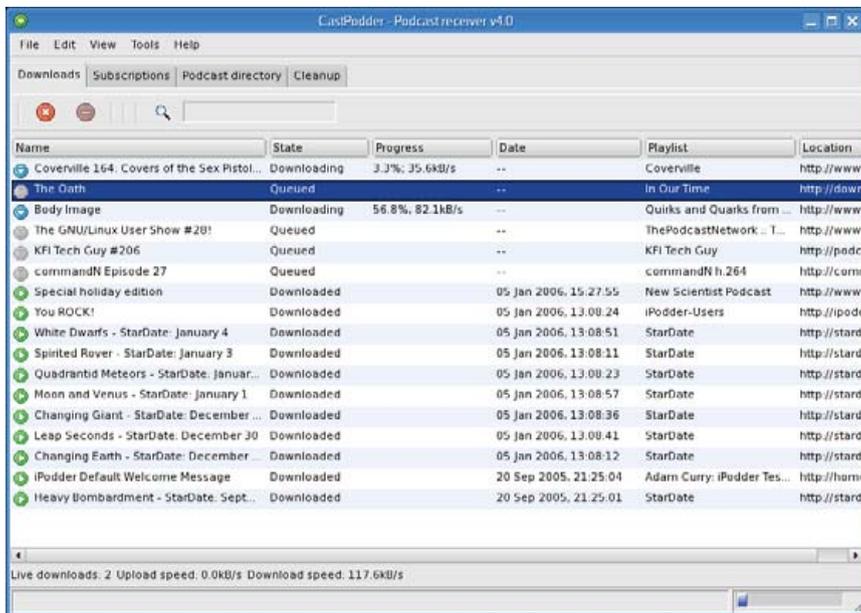


Figure 4. The Downloads tab tells you the progress of each podcast download.

and its location on the Web. Look to the far left next to the episode title, and you'll see a check box. Click this box to flag the episode for download. If there is a little green arrow next to the episode title, that means it already has been downloaded to your Linux system.

Once you have selected the episodes you want, click the first button on the left of the icon bar, the one that says Scan for new Feeds when you pause your mouse cursor over it. The button will turn gray, and the download of your podcasts will begin. To see the progress of those downloads, click on the Downloads tab. You'll see the current state of the download, as well as the speed and percentage completed of the transfer (Figure 4).

Finally, we get to the whole point of finding and downloading all these shows, and that's listening to them. Take a look again at that little green arrow to the left of your downloaded episodes (this applies to both the Download and the Subscriptions tab). Either double-click the entry or right-click and select Play episode in media player from the pop-up menu. By default, XMMS will start to play the program you've chosen (Figure 5). It's possible to override this choice of player, but your choices are limited to XMMS, the Beep Media Player and no player at all. To make the change, click File on the menu bar and select Preferences. Then, in the Preferences window, click on the Player tab.

CastPodder swallows down into a nice, lemony applet in your system tray so that the program can keep running out of the way. This is particularly handy when you consider CastPodder's scheduling functionality. From the Subscriptions tab, click the Scheduler button on the icon bar to automate your podcast downloads. Automatic updates and downloads to subscribed feeds can be set to run at a specific time of day or at regular intervals



Figure 5. By default, CastPodder launches XMMS to play your podcasts.

throughout the course of the day. That way, you don't have to spend time checking up on new episodes. Just let CastPodder do the work for you.

It appears, *mes amis*, that closing time is almost upon us. Still, I should mention one final CastPodder feature before I finish and that has to do with cleaning up. Earlier on, I mentioned that it is good to have lots of free space if you get into the habit of downloading several podcasts. It is also a good idea to clean up those old shows as new ones continue to be downloaded. Click the Cleanup tab and CastPodder lets you select individual episodes by the podcast feed name and easily delete them. It is also possible to select them all with a single click should you want every episode deleted.

It is now closing time, but I see that many of you are in the middle of searching for and subscribing to many interesting programs in the podcastverse. Take a little more time, *mes amis*. I'm sure that François will be more than happy to refill your glasses one final time before we say, "*Au revoir*". Hmm...perhaps one of you can check to see if there are any podcasts dedicated to wine. 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/8829. ■

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

NAS — LINUX BASED. PERFORMANCE TUNED.

You chose a Linux environment for many reasons, not the least of which includes its performance, reliability, security and open environment. Why not choose Network Attached Storage (NAS) that meets the same criteria? Introducing the Aberdeen AberNAS 170 and AberNAS 270 series of Linux based NAS appliances. They deliver all the features you love about Linux, in the highest performing, most competitively priced NAS in the marketplace today. Just look at these features:

- Best fit for Unix/Sun Micro/Linux environments or mix of Windows and non-Windows environments
- Browser based GUI management
- Substantially faster NFS and CIFS performance compared to Windows Storage Server
- Extremely virus resistant—Disc-On-Module (DOM) operating system provides integrity, avoids corruption and allows for easy upgrade
- Embedded iSCSI target with unlimited initiators
- Schedulable volume snapshots
- PCBackup Utility included allows multiple file versions and instant backup of any modified file
- NAS-2-NAS Replicator included
- Available with 500GB 8.5ms 7200rpm hot-swap SATA II 3Gb/s hard drives with 16MB cache featuring Rotational Vibration Safeguard and staggered (delayed) spin-up
- SATA II 3Gb/s enables the highest level of performance at low cost of ownership
- 300MBytes/sec per drive data transfer rates—twice as fast as regular SATA and three times as fast as Ultra ATA/100
- Single controller for up to 24 drives—up to single 16TB partition
- Intel RAID 6 Engine to support extreme performance RAID 6—can protect mission-critical data from two concurrent disk drive failures
- Status indication via HDD activity/fault LEDs and audible alarm
- Management port/firmware supports browser-based RAID manager, SMTP email notification, SNMP agent
- Native Command Queuing (NCQ) for extreme performance
- "Smooth Stream" technology for enhanced video streaming
- Apple support including AFP 3.0
- **1TB AberNAS solutions starting at \$2,345.**



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

CRN Test Center Recommended August 22, 2005—Aberdeen Stachoven A381

Finalist: Best Servers

LinuxWorld Fraud & Excellence Awards August 8-11, 2005—Aberdeen Stachoven A261

"solid redundancy for data drives, OS, and power"

FCMagazine April 26, 2005—Aberdeen AberNAS 211

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

FCMagazine February 8, 2005—Aberdeen Stachoven A261S

"extremely approachable and easy to use at a very affordable price."

CRN Test Center Recommended January 10, 2005—Aberdeen AberNAS i100

"terrific for video serving or other storage-intensive tasks"

FCMagazine October 5, 2004—Aberdeen XIDAS



ABERDEEN ABERNAS 170 SERIES

1U NAS with Hardware RAID

- Up to 2TB of Network Attached Storage in a 1U format
- Intel® Pentium® 4 processor 531 supporting Hyper-Threading Technology, 3.0E GHz with 800MHz FSB and 1MB Cache
- 1GB Dual Channel DDR400 SDRAM (Low Latency)
- Single controller for all drives with RAID level 0, 1 (10), 3 and 5
- 4 x Hot-Swap Hard Drives with single partition capability
- Slim DVD Drive
- 300W Power Supply
- Linux based proprietary NAS software (see above for features and benefits)
- Disc-On-Memory (DOM) operating system provides integrity and avoids corruption
- **5-Year Limited Warranty**

AberNAS 170 **\$3,195**

- 1TB of Network Attached Storage
- 4 x 250GB Hot-Swap SATA Hard Drives with 8MB cache

AberNAS 176 **\$3,895**

- 1.6TB of Network Attached Storage
- 4 x 400GB Hot-Swap SATA Hard Drives with 8MB cache

AberNAS 178 **\$4,495**

- 2TB of Network Attached Storage
- 4 x 500GB Hot-Swap SATA II 3Gb/s Hard Drives with 16MB cache

ABERDEEN ABERNAS 270 SERIES

2U NAS with Hardware RAID

- Up to 4TB of Network Attached Storage in a 1U format
- Intel® Pentium® 4 processor 531 supporting Hyper-Threading Technology, 3.0E GHz with 800MHz FSB and 1MB Cache
- 1GB Dual Channel DDR400 SDRAM (Low Latency)
- Single controller for all drives with RAID level 0, 1 (10), 3, 5 and 6
- 8 x Hot-Swap Hard Drives with single partition capability
- Slim DVD Drive
- 460W Redundant Power Supply
- Linux based proprietary NAS software (see above for features and benefits)
- Disc-On-Memory (DOM) operating system provides integrity and avoids corruption
- **5-Year Limited Warranty**

AberNAS 271 **\$4,995**

- 2TB of Network Attached Storage
- 8 x 250GB Hot-Swap SATA Hard Drives with 8MB cache

AberNAS 272 **\$6,995**

- 3.2TB of Network Attached Storage
- 8 x 400GB Hot-Swap SATA Hard Drives with 8MB cache

AberNAS 273 **\$7,995**

- 4TB of Network Attached Storage
- 8 x 500GB Hot-Swap SATA II 3Gb/s Hard Drives with 16MB cache



DAVE TAYLOR

Calling All Functions, and Some Math Too!

Shuffle off to the next step in creating a shell script *Blackjack* game.

If you've been keeping track of my column, you'll know that we're building a *Blackjack* game as a shell script. Why? Because most shell scripts are far too boring to study without nodding off, so considering all the nuances of shell script programming within the context of a game just sounds more fun!

My last column talked about how to shuffle a deck of cards, as portrayed in a program with a simple array of 52 values, 1–52. There are some interesting nuances to the shuffle problem. Let's start there. Then we'll look at how to turn an arbitrary 1–52 value into a familiar rank and suite from a deck of cards.

The Shuffle Function

If you've been writing shell scripts since the dawn of UNIX, you might not have realized that modern shells now support functions and procedures, just like a "real" programming language. For any block of code that you plan on executing more than once, it's the way to go.

Here's the shuffle code, written as a shell function:

```
function shuffleDeck
{
    count=1

    while [ $count -le 52 ]
    do
        pickCard
        newdeck[$count]=$picked
        count=$(( $count + 1 ))
    done
}
```

This builds an array called `newdeck`, which is actually the shuffled deck (in the last column we showed `deck`, which was an array of cards in linear order), and it uses some basic shell math with the `$(())` notation to save spawning a subshell to invoke `expr` for each increment of the count variable.

I said that shell scripts are robust programming environments and that might be a tiny bit of hyperbole, really. Eagle-eyed readers will notice that the `pickCard` function returns its value by setting a global variable, `picked`, which isn't really optimal programming strategy. But it works, and pragmatism is an important part of any good approach to software development, isn't it?

The full `pickCard` function is key to making this work, but it is too long to include here, so grab it from the *LJ* FTP site for your reading pleasure (ftp.ssc.com/pub/lj/listings/8774.tgz).

With the `shuffleDeck` function written and an `initializeDeck` function, as shown here:

```
function initializeDeck
{
    card=1 while [ $card -le 52 ] do
        deck[$card]=$card card=$(( $card + 1 ))
    done
}
```

it's easy to do the rudiments of shuffling the deck and dealing out two cards for the player and two for the dealer:

```
initializeDeck shuffleDeck

echo "** Player's hand: ${newdeck[1]}, ${newdeck[3]}"
echo "** Dealer's hand: ${newdeck[2]}, ${newdeck[4]}"
```

Let's run this and see what kind of results we get:

```
$ ./blackjack.sh
** Player's hand: 22, 49
** Dealer's hand: 11, 8
$ ./blackjack.sh
** Player's hand: 19, 32
** Dealer's hand: 49, 10
$ ./blackjack.sh
** Player's hand: 44, 23
** Dealer's hand: 46, 11
```

Displaying cards as a numeric value from 1–52 is not the most friendly, so let's turn our attention to the display of the card values in the familiar rank and suite of a traditional deck of playing cards.

Math Games to Identify Rank and Suite

A deck of cards is composed of 52 cards, split evenly into four suites of 13 cards. The order of the suites doesn't matter (in *Blackjack*, at least), but the rank does. Indeed, the goal of the game is have a summary rank value of 21 points without going any higher.

The rank of a card is the remainder of the numeric card value divided by 13. In math terms, this is called the modulus and can be computed thusly:

```
rank = cardvalue % 13
```

Reclaim lost time



Levanta Intrepid™

The world's first Linux management appliance

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

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

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

30-Day
Money-Back Guarantee

**Visit Us at LinuxWorld
Boston - April 3-6, 2006**


LEVANTA
www.levanta.com
1.877.LEVANTA



To put this into proper shell notation, we'll again use the `$(())` shortcuts and end up with:

```
rank=$(( $card % 13 ))
```

Getting the suite should be simple too; it's the other half of the division. In other words, if the card value is 17, then $17/13 = 1$, meaning it's suite #1, and $17\%13 = 4$. Since we'd like to have our suites in the range of 1–4, rather than 0–3, however, we'll need to add one to the equation. Further, the 13th of each card is the same suite as the earlier 12, so we'll also have to subtract one before we do the division (card #13 would be suite #1 if we just calculated $13/13$, but if we calculate $12/13$ for that, we'll correctly identify it as part of suite #0).

This is very confusing, so here's the equation instead:

```
suite=$(( ( ( $card - 1 ) / 13 ) + 1 ))
```

Much clearer, right? Seriously, you can experimentally verify that this works correctly. The important edge cases are `value=1`, `value=12`, `value=13` and `value=14`. If you can get those right, you're good for all values in the deck.

Once we've identified the rank and suite of a card, we

just have to do a bit of fancy footwork to turn numbers into words:

```
case $suite in
  1 ) suite="Hearts" ;;
  2 ) suite="Clubs"  ;;
  3 ) suite="Spades" ;;
  4 ) suite="Diamonds" ;;
  * ) echo "Bad suite value: $suite"; exit 1
esac
```

and:

```
case $rank in
  0 ) rank="King" ;;
  1 ) rank="Ace"  ;;
  11) rank="Jack" ;;
  12) rank="Queen" ;;
esac
```

Put these together in a function called `showCard` (which returns `$cardname` as the calculated rank and suite), and we can now clean up a bit:

```
initializedDeck shuffleDeck
```

```
echo -n "** Player's hand: "
showCard ${newdeck[1]} ; echo -n "$cardname, "
showCard ${newdeck[3]} ; echo "$cardname"
```

```
echo -n "** Dealer's hand: "
showCard ${newdeck[2]} ; echo -n "$cardname, "
showCard ${newdeck[4]} ; echo "$cardname"
```

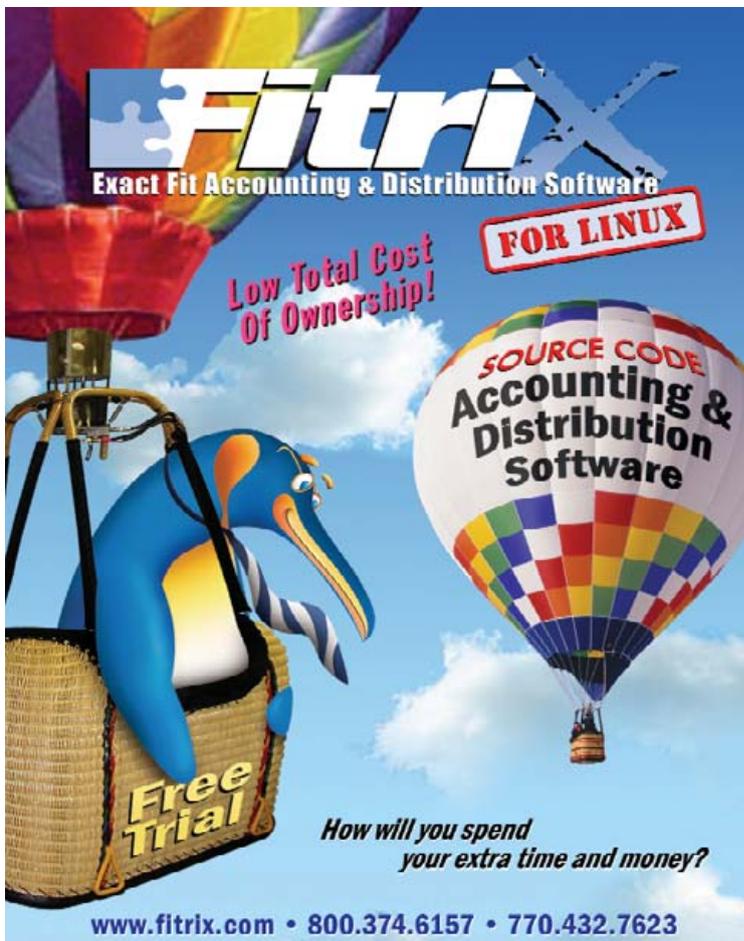
Now we can start to see the game come together! Consider:

```
$ ./blackjack.sh
** Player's hand: 8 of Clubs, 3 of Diamonds
** Dealer's hand: King of Spades, 3 of Spades
$ ./blackjack.sh
** Player's hand: 2 of Spades, 4 of Spades
** Dealer's hand: 10 of Spades, 4 of Hearts
```

Let's stop here this month as that's a lot of code to dig through already. I invite you to pop over to the *LJ* FTP site to grab all the source code so far, so you can experiment with this script yourself too.

Next month, we'll start looking at the game logic itself, but for now, Vegas beckons for a big tradeshow and, well, I can write it off as research for *Linux Journal*, can't I? ■

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.



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® Optron™ Dual-Core Processors 200+ with Hyper Transport Technology
- Upto 64GB 400MHz ECC DDR (16 Sockets)
- Upto 2TB 4 x 500G Swap Drive per 1U Rack
- 1 x PCI-X 133MHz, 1 x PCI-E Slot
- Dual Gigabit Ethernet, ATI Graphics, 4 SATA RAID
- Slim CD-ROM Drive, Optional DVD-RW or CDRW
- Optional Slim Floppy Drive
- 1U 24" Depth Rack Chassis with upto 600W P/S
- 4 x Swappable Drive Bays (SATA or SCSI)
- Supports Linux, FreeBSD or Windows
- Custom Configuration Available
- Please call for other Options



1114AIS-2050M, 280+, 32GB, 2TB **\$12,950**
 1114AIS-2050M, 265+, 16GB, 1TB **\$6,299**
 1112ES-2200A, 244+, 2GB, 500G **\$1,999**



4024AIS 4U 12TB NAS/SAN **\$14,999**
 4024AIS 4U 6TB NAS Storage **\$7,999**

Linux Appliance PCs

- Custom Made Odd Size Chassis
- AC or DC Power Supply
- Low Power Voltage AMD Sempron™ Processor or High Performance AMD Athlon™ G4 Processor
- Diskless or Flash OS Boot Drive
- Swappable Hard Drive, CD-ROM, FDD
- Integrated Graphics, Ethernet, USB
- Optional LCD LED Control Module
- IS2 Audio, MPEG2/4 Hardware Video
- Upto 4 Ethernet Ports or 4 Serial Ports
- We have over 18 years OEM Experience in Set top Box, Digital Media Player, POS Kios, Thin Client, Networking Appliance.



SX2500SPLJ11B
 OEM Appliance starts at **\$299**

64GB DDR 4-way Workstation

- 2 x AMD® Optron™ Dual-Core Processors 280+
- Upto 64GB 400MHz ECC DDR (16 Sockets)
- 2 x 133MHz PCI-X, 1 x PCI-E x16 Slots
- Dual Gigabit Ethernet, 4 x SATA-RAID Controller
- Quiet and Cool 12 Bay lower + 600W 80Plus P/S
- 250G HD, DVD-RW, Floppy, Optional Card Reader
- 8-Layer Motherboard with Special Quiet Cooling
- On-board ATI Graphics, Sound Blaster 7.1 Sound
- Supports 64/32 bit Linux, FreeBSD or Windows
- Special for Large Memory Intensive Applications
- Built-to-Order or Configure-to-Order



64G RAM, 2x280+, QuadroFX1500 **\$22,999**
 32G RAM, 2x265+, QuadroFX3400 **\$9,999**
 16G RAM, 2x244+, QuadroFX1300 **\$4,699**

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

1U Power Saving ISP Server

- AMD Sempron™ or Optron™ Processor
- 512M DDR 400MHz Memory
- 80GB Hard Drive
- 10/100Mbit Ethernet
- 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® Optron™ Dual-Core Processors 865+ with Hyper Transport Technology
- Upto 128GB DDR Memory for 16-way (32 sockets)
- Upto 64GB DDR Memory for 8-way (16 sockets)
- 4x Gigabit LAN, 8x SATA RAID-5 for 16-way
- 2x Gigabit LAN, 4x SATA, U320 SCSI for 8-way
- 4 x 133/100/66MHz PCI-X Slots for 16-way
- 2 x 133, 2 x 66MHz PCI-X, 1x PCI Slots for 8-way
- On-board ATI Graphics, USB 2.0
- 5U 26" Rack 1300W 3+1 Redundant P/S 16 way
- 2U 27" Rack 700W PFC P/S for 8-way
- Supports Linux, FreeBSD or Windows
- Custom Configuration Available
- Please call for other Options



5U 16 way 865+, 32GB, 2TB, 8800U5 **\$25,999**
 2U 8 way 865+, 8GB, 750GB, 8422C **\$9,500**

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

939NX4LX 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



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

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



MICK BAUER

Security Features in SUSE 10.0

SUSE is a security-friendly distribution with a plethora of security-related tools.

Over the years, we've seen more and better security features incorporated into our favorite Linux distributions. Distribution-specific security awareness manifests itself in many ways, including:

- Availability of security-enhancing applications.
- "Hardening" functionality in setup/installation scripts.
- The way patches are handled.
- Default settings of network applications.

This month, I begin a series of three articles on distribution-specific security in SUSE Linux, Debian GNU/Linux and Red Hat Enterprise Linux. These are the three distributions with which I've had the most experience, and they are arguably the three most popular. (But as with anything, if you want to contribute an article about your own favorite distribution, go for it! See our author's guide at www.linuxjournal.com/xstatic/author/authguide.)

I'll start with SUSE 10.0. SUSE is a general-purpose, commercially produced Linux distribution developed for Intel 32- and 64-bit platforms. Originally based in Germany and still primarily developed there, SUSE is now owned by Novell. There are a number of different SUSE products, including SUSE Linux, a "personal" version available from numerous retail outlets; SUSE Linux Enterprise Server, an "enterprise-grade" version available directly from Novell; and OpenSUSE, which is essentially the same as SUSE Linux but without installation media (it's installable only over the Internet), printed manuals or installation support.

The basis of this article is SUSE Linux 10.0, that is, the commercial "personal use" version. Everything I say here should be equally applicable to OpenLinux 10.0, and mostly relevant to the Enterprise versions of SUSE. Presumably, the Enterprise versions include additional security-related packages and features.

Installing SUSE Linux 10.0

System security begins with installation. This is your first opportunity to make crucial decisions concerning what role the system will play, which software the system will run and how the system will be configured. Therefore, it's useful to begin our discussion of SUSE security with the installation process.

All versions of SUSE use YaST (Yet Another Setup Tool) both for initial system installation and for ongoing system administration. Over the years, YaST has evolved from a simple RPM front end to a modular, comprehensive administration tool that can be used to configure not only low-level system software but also complex server applications such as Apache and Postfix.

We'll talk more about YaST shortly. Your immediate problem during initial OS installation, however, is deciding which software packages to install. And if you're security-focused, this is a happy problem. SUSE Linux 10.0 offers a wide variety of security applications from which to choose.

In my view, these applications fall into two categories: system security applications and security-scanning applications. The former include both general-purpose applications with strong security features—Postfix springs instantly to mind—and applications whose sole purpose is providing security controls to other applications or to the underlying operating system, of which `tcpwrappers` is a classic example. Table 1 lists the packages in SUSE Linux 10.0 that enhance system security.

Actually, the lengthy list of packages in Table 1 represents only particular favorites of mine and SUSE-specific selections. SUSE includes many, many more system security tools, including `tcpd` (`tcpwrappers`), `openssl`, `chkrootkit`, `sudo` and `wipe`. You can view the full list of packages included in SUSE Linux 10.0 at www.novell.com/products/linuxpackages/professional/index_all.html.

Besides securing the system on which you install SUSE, you may be interested in using a SUSE system to validate the security of other systems or of entire networks. SUSE is a good choice for this. Table 2 shows some SUSE Linux 10.0 packages that can be used for security scanning. Note that you should never install these packages (except perhaps `Snort`) on any Internet-connected server. Each is of much greater use to an attacker than it is to you in that context. Scanning software should be performed from systems that are normally kept out of harm's way.

If you're new to SUSE, you should be aware that by default, YaST uses a Selections filter (*view*) for selecting packages, in which only a small subset of all available packages is offered to you. If you don't see something you need in this view, for example, `nessus-core`, use the Package Groups filter to see a more complete set of categories. If you want to see a single list of all packages in alphabetical order, simply set the filter to Package Groups and click on the group `zzz All` (Figure 1).

You also can set the filter to Search to search for packages by name or keyword.

After you've selected and installed all software packages, YaST allows you to set the root password and create the first (nonroot) user account. By default, SUSE uses Blowfish for password encryption, and YaST checks the password you type for complexity. (Too-simple a password can be easily guessed or brute-force cracked by an attacker.)

You're also given the opportunity to enable local firewall scripts (enabled by default), and the SSH and VNC remote-shell daemons (both disabled by default). Note that of the latter two,



~~OUR CRM SYSTEM IS SO INFLEXIBLE.~~
~~OUR CRM SYSTEM COSTS TOO MUCH.~~

Frustrated by CRM restrictions?

Take back control of your CRM solution with SugarCRM, the most powerful commercial open source CRM solution for managing sales, marketing, support, and group collaboration activities.

Learn More: www.sugarcrm.com/linuxjournal

Complimentary white paper, *"Running a Business on Commercial Open Source."*

 **LINUXWORLD**
CONFERENCE & EXPO

Boston, USA • April 3 - 6, 2006 • Booth 1100

See **SugarCRM** demos
and speak to CRM experts.

FREE CRM Solution Pack CD

Commercial Open Source CRM Res

www.sugarcrm.com

Commercial Open Source Customer Relationship Management

 **SUGARCRM**™

Copyright © 2006 SugarCRM, Inc. All rights reserved. SugarCRM and the SugarCRM logo are trademarks of SugarCRM, Inc. in the United States, the European Union and other countries.

www.sugarcrm.com • www.sugarforge.org
1.875.SUGARCRM • +1.408.454.6941

Table 1. Some Security-Enhancing Packages in SUSE Linux 10.0

Package Name	Description
aide, fam	File integrity checkers, both similar to Tripwire.
bind-chrootenv	Automatically creates a chroot environment in which to run BIND (the DNS daemon) more securely.
clamav, antivir	Antivirus packages—clamav is completely free, but antivir is commercial (free for personal use).
cracklib	Library and utilities to prevent users from choosing easily guessed passwords.
gpg, gpg2, gpa	GNU Privacy Guard (gpg), a versatile and ubiquitous e-mail- and file-encryption utility.
ipsectools, openswan	Tools for building IPsec-based virtual private networks.
openldap, freeradius	Open-source authentication daemons.
proxy-suite	An FTP security proxy developed by SUSE.
seccheck	SUSE-customized cron scripts that perform various security checks against logs, system state and so on, and send e-mail reports to you.
subdomain-utils, subdomain-profiles, mod-change-hat and so on	AppArmor, a mandatory access control (MAC) system that restricts the behavior of specific binaries. SUSE uses this instead of SELinux, which it closely resembles.
squid, SquidGuard	Squid is a popular HTTP/HTTPS proxy. SquidGuard adds access controls and other security features.
SUSEfirewall	SUSE's handy front end for Linux's netfilter/iptables.
syslog-ng	Advanced system logger, much more powerful than syslogd. syslog-ng is SUSE's default logger.
tinyca2	Front end to OpenSSL for managing Certificate Authorities.
yast2-firewall	Firewall functionality.
vsftpd	The Very Secure FTP Daemon.
xen, FAUmachine, uml-utilities, bochs	The Xen, FAUmachine, User Mode Linux and BOCHS virtual machine environments.

Table 2. Security Scanners in SUSE Linux 10.0

Package Name	Description
ethereal, tcpdump	Excellent packet sniffers.
fping	Flood ping (multiple-target ping).
john	John the Ripper, a password-cracking tool (legitimately used for identifying weak passwords).
kismet	Wireless LAN sniffer.
nessus-core, nessus-libraries	The Nessus general-purpose security scanner.
nmap	Undisputed king of port scanners.
snort	Outstanding packet sniffer, packet logger and intrusion detection system.

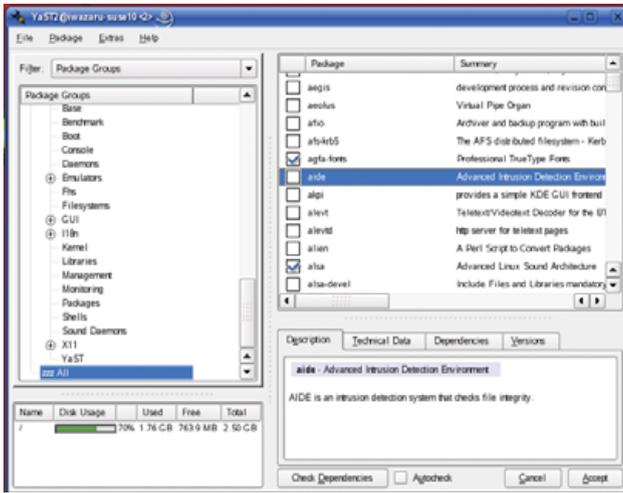


Figure 1. Viewing All Available Packages in YaST

SSH is the best choice for administering bastion hosts (hardened Internet servers)—among other reasons, you shouldn't be using the X Window System on bastion hosts unless you've got a very specific, very compelling reason. YaST, it should be

noted, runs perfectly well in text (ncurses) mode, with exactly the same modules and options as the X version. Also, tightvnc, the version of the VNC remote-desktop tool shipped with SUSE, doesn't encrypt session data, only authentication data.

Note also that at installation time, you aren't given the opportunity to customize your local firewall settings. Initially, a default script is used that provides a simple "allow all outbound transactions, allow nothing inbound that wasn't initiated locally" policy. In other words, the default SUSEfirewall script is perfectly appropriate for most desktop systems, but it is inadequate for server use. You can change this later on by running YaST's Firewall module.

YaST then lets you choose from the following methods for authenticating nonroot users:

- local /etc/passwd file (default).
- LDAP.
- NIS.
- Samba (Windows NT Domains).

Active Directory authentication is also supported in SUSE

Hurricane Electric Internet Services...**Speed and Reliability** That Lets You Soar Over the Competition!

Flat Rate
Gigabit Ethernet

1,000 Mbps of IP

\$13,000/month*

Full 100 Mbps
Port

Full Duplex

\$2,000/month

Colocation Full
Cabinet

Holds up to 42 1U
servers

\$400/month

Order Today!

email sales@he.net or call 510.580.4190



* 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 London, UK; NIKHEF in Amsterdam, NL; Hurricane in Fremont, CA and Hurricane in San Jose, CA

Linux 10.0, via Kerberos.

Once you've selected an authentication method, you can create your first nonroot user account. Be sure to leave Automatic Logon disabled unless your system has very low security requirements indeed—enabling this causes the machine to log in your nonroot user automatically at boot time. (About the only situation in which this is a good idea, I think, is for kiosk-type systems!)

And that's it—SUSE installation is now finished! Your job as a security-conscious system administrator, however, is not.

Security-Related YaST Modules

After the first time you boot your newly minted SUSE Linux system, you immediately should log in as your unprivileged user and

invoke YaST. If you do this from within KDE or GNOME, you'll be prompted for the root password automatically, but in a text-console session, you need to use `su -c` to invoke `/sbin/yast`.

As I mentioned earlier, YaST has a lot of security functionality built in. YaST modules particularly relevant to system security are listed in Table 3.

Of these YaST modules, Online Update is one of the most important. You immediately should use it to configure automatic patch downloads and, unless your system is under a change-control process, automatic patch installation as well. YaST Online Update was one of the first automatic patch utilities offered in a major Linux distribution, and it's still one of the best. Use it to take advantage of SUSE's excellent record of providing prompt, well-tested security patches.

Table 3. Security-Related YaST Modules

YaST Section	Module Name	Description
Software	Online Update	Sets up manual and automatic software updates.
	Software Management	For installing and removing packages.
	Virtual Machine Installation (XEN)	Creates virtual machines for the Xen 3 virtual machine environment.
System	/etc/sysconfig Editor	Edits <code>dæmon</code> startup parameters.
	System Services (Runlevel)	Manages startup scripts.
	Powertweak	Sets advanced kernel parameters, such as TCP timewait sockets.
Network Services	DNS Server	Configures BIND.
	HTTP Server	Configures Apache.
	LDAP Client	Sets up LDAP authentication and lookups.
	Mail Transfer Agent	Configures Postfix or Sendmail.
	Kerberos Client	Sets up Kerberos authentication, including Active Directory.
	Remote Administration	Configures TightVNC.
Novell AppArmor	Various	For managing AppArmor mandatory access controls on specific binaries.
Security and Users	Firewall	For managing netfilter/iptables settings.
	Local Security	Determines password complexity and length, password aging, file-permission schemes and various other system security parameters.
	Group Management	Used to create, edit and delete group accounts.
	User Management	Used to create, edit and delete user accounts (actually the same module as Group Management, which is dual-purpose).

BOSS TTON



USENIX
Annual
Technical
Conference '06

May 30–June 3, 2006

Join leading researchers for 5 full days on the latest in cutting-edge technology

- **5 DAYS OF TRAINING**
 - Tuesday–Saturday, May 30–June 3
 - 34 full- and half-day tutorials
 - Learn from expert instructors
 - Topics range from Ajax to VoIP

- **3-DAY TECHNICAL PROGRAM**
 - Thursday–Saturday, June 1–3
 - Systems Practice & Experience Refereed Papers Track
 - Keynote by Larry Peterson, Princeton and Planetlab
 - Plenary by Greg Brandeau, Pixar
 - Closing by Trevor Blackwell, Anybots

Plus invited talks, guru sessions, BoFs, poster sessions, receptions, and more!

Register by May 12 and save up to \$300 • www.usenix.org/usenix06/lj

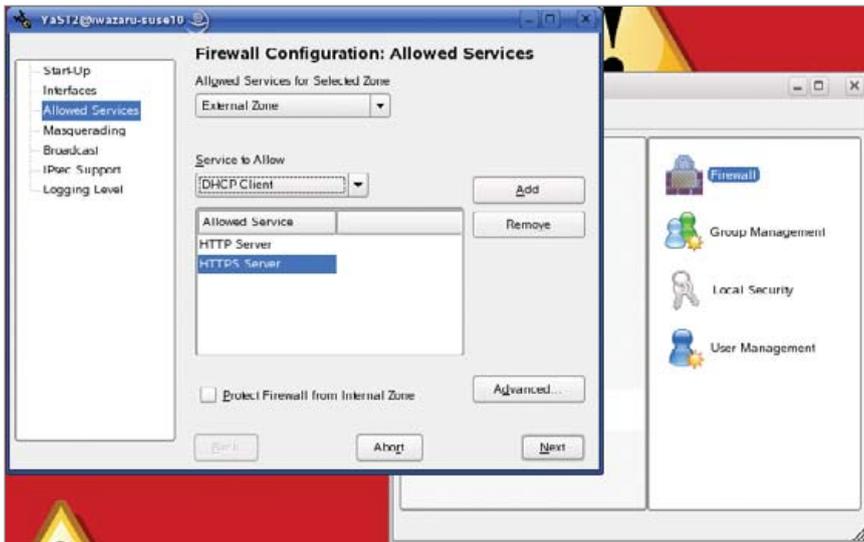


Figure 2. YaST's Firewall Module

The Firewall module (Figure 2) is also extremely useful, especially if you're uncomfortable creating and managing your own firewall scripts (I acknowledge that people like me, who find this fascinating and fun, are rare). Similarly, Group/User Management eliminates the need for you ever to edit `/etc/group` or `/etc/passwd` manually.

The Virtual Machine Installation module and Novell AppArmor section are also especially noteworthy. So much so, in fact, that I should spend some time talking about SUSE's virtual machine and mandatory access control systems, respectively, in a little more depth.

Virtual Machines in SUSE Linux

You may recall my article "The Future of Linux Security" [LJ, August 2005], in which I touted virtual machine environments and hypervisors (aka security monitors) as being an important new direction in system security. If you don't recall this, the gist of it is that it's because MAC schemes such as SELinux are viewed by many people as too complex. A simpler approach instead is to run each major application or service on its own virtual machine. That way, if for example a virtual machine in which Sendmail is running gets compromised, a virtual machine running Apache2 on the same physical hardware won't be in immediate or direct danger.

Virtual machines, therefore, provide a powerful and easy-to-understand means of isolating complex applications from each other. And, SUSE Linux 10.0 includes no fewer than three different virtual machine technologies.

The Xen 3 environment, which originated at Cambridge University, is provided by SUSE as a "technology preview". To the best of my determination, this simply means that because Xen 3 is an immature and potentially unstable application, SUSE is simply trying to lower people's expectations of its usability—the version of Xen 3 in SUSE Linux 10.0 isn't a special preview or evaluation version or anything like that. Xen 3 supports Linux, FreeBSD, NetBSD and Plan9 "guest" (virtual) systems.

Alternatively, the FAUmachine virtualization environment includes RPM packages that enable support for SUSE 9,

Debian 3.0, OpenBSD 3.5/3.6 and Red Hat 9 guest systems. One advantage of FAUmachine over Xen 3 is that in FAUmachine, the guest systems' kernels run on the host system with nonroot (unprivileged-user) permissions.

User Mode Linux is another virtualization environment offered in SUSE Linux 10.0 via the `uml-utilities` package. Like FAUmachine, its guest kernels run without root privileges.

Novell (Immunix) AppArmor

However, not everyone has given up on MAC-based system security, and SUSE has covered this area handsomely by acquiring and repackaging Immunix's AppArmor (aka Subdomain). AppArmor is similar to SELinux, in that it allows you to restrict the behavior of specific processes, with an effect similar to but more effective than running them in chroot jails.

(Note that although SUSE provides the `libseline` package and includes SELinux functionality in its default kernel, SELinux isn't officially supported in SUSE Linux. You need the packages available at www.cip.ifi.lmu.de/~bleher/selinux to run SELinux in SUSE Linux.)

The document `/usr/share/doc/packages/subdomain-docs/ug_apparmor.pdf`, included in the `subdomain-docs` package, is the AppArmor User's Guide, and it tells you everything you need to know about configuring and using AppArmor. Suffice it to say for now that if you simply run the YaST AppArmor Control Panel module and enable AppArmor, a default profile is loaded that includes settings for many common daemons and commands, including `netstat`, `ping`, `traceroute`, `firefox`, `evolution`, `gaim`, `syslogd`, `acoread`, `ethereal`, `appropos`, `procmail`, `postfix` (`smtpd`, and so on), `Apache2` (`httpd2-prefork`), `nscd`, `identd`, `ntpd`, `sshd` and `squid`.

This is a limited-feature version of AppArmor, so apparently it provides only a subset of features available in the full \$1,250 US version. Personally, I'm not clear as to precisely what the difference is, though—everything I tried to do with the version in SUSE Linux 10.0 seemed to work fine, so this would not appear to be a too significantly crippled edition. Perhaps the full version includes a longer list of preconfigured applications.

Conclusion

These aren't SUSE Linux 10.0's only security features. I haven't talked about how secure many applications' default settings are (in general they're quite secure, with daemons running with nonroot privileges whenever possible, network listeners such as `sshd` typically disabled by default and so on).

This is a very security-friendly version of SUSE Linux indeed. Remember, though, that real security begins with you—little of SUSE's security potential is realized until you configure or at least enable it yourself! Hopefully, this article has helped you get a feel for what that potential is.

Next month, it's on to Debian 3.1. Until then, be safe! ■

Mick Bauer (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".



the POWER of JAVA™



Don't Miss the 2006 JavaOne™ Conference

More than 300 technical sessions and Birds-of-a-Feather sessions will be offered at the 11th Annual JavaOne™ Conference. See and hear from the Industry leaders and technology experts over four content-rich days. Included are:*

PLATINUM COSPONSORS



GOLD COSPONSORS



SILVER COSPONSORS



Introduction to AJAX

Ben Galbraith, *Consultant* | Dion Almaer, *Adigio, Inc*

This session provides an introduction to AJAX and an orientation to the state of the AJAXian universe. It demonstrates the basic AJAXian techniques through live coding and demonstrates and deconstructs more-advanced examples of AJAX.

Effective Java™ Reloaded

Joshua Bloch, *Google, Inc.*

It has been five years since *Effective Java™* was released. The Java platform has evolved, and we've learned more about how to use it to best effect. Therefore, a second edition of *Effective Java™* is being released to coincide with the 2006 JavaOne conference. This presentation covers new material that has been added to the second edition, material that should be useful to every working Java technology programmer.

Spring Framework Update

Rod Johnson, *Interface21*

Rod Johnson, the father of Spring, brings attendees up to date on some of the many significant new features in the Spring 1.3 and 1.4 releases. He discusses what's new and cool in the Spring world and examines the implications of these new features for best practice in developing applications with the Spring Framework. Johnson shows code examples throughout the presentation, leaving attendees ready to try these features out for themselves.

To see more information on the conference offerings visit java.sun.com/javaone/sf

*Content subject to change.

Register by April 14, 2006 and SAVE \$100.
java.sun.com/javaone/sf



JavaOne™ Conference | May 16-19, 2006
JavaOne™ Pavilion: May 16-18, 2006, Moscone Center, San Francisco, CA

JavaOne™
Sun's 2006 Worldwide Java Developer Conference



DOC SEARLS

The Producer Electronics Revolution

The open media revolution is upon us with blogs and podcasts leading the way.

I'm writing this on the way back from the Consumer Electronics Show (CES) in Las Vegas—the biggest tradeshow on Earth. This year, 150,000 visitors came to see 2,500 exhibits spread across 1.3 million square feet. One of those visitors was Andrew Leyden of PenguinRadio, who joined me for a day of hunting for Cool Linux Stuff. Andrew is a consumer electronics veteran, making Linux-based Internet radios since 1999.

Andrew is a walking case study in resourcefulness and persistence, as well as living proof that, as he puts it, “The commercial sphere is shrinking and the DIY sphere is growing.”

You wouldn't have heard a statement like that from the brand-name companies and keynoters who starred at CES. Not from Bill Gates, who gave the opening keynote (which, in a long-standing tradition for me, I missed—though thousands waited hours to attend). Not from Intel, which brought Tom Hanks, Danny DeVito, Morgan Freeman and other stars on stage to promote its proprietary new “content delivery” system. Not even from Larry Page of Google, who gave a sincere, human and blessedly unscripted keynote, guest-starring Robin Williams, who was funny as hell.

CES is a commercial sphere, and most of us continue to live in commercial environments. That's what proprietary systems are. And, that's why Google (which runs its massive search infrastructure on Linux and participates in many open-source development communities) introduced a raft of products and services at CES that ran only on Windows clients. One was a new video store that worked only for Windows clients built by Google, and which required Google's own DRM. When I asked Larry if and when Google would come out with stuff that ran on other clients, he admitted that it was “a problem” and said they're working on it. At least when Google says that, I believe them. Out on the CES show floor, when an executive with another company answered the same kind of question with “We're always looking to improve the user experience”, Andrew and I could barely conceal our utter disbelief.

Of course, everybody talked about “putting the user in charge” and “having your media your way”, but there was little credit given to users who really do take charge, operate independently and even produce their own media. In his 1995 essay “Death From Above”, John Perry Barlow wrote, “America remains a place where companies produce and consumers consume in an economic relationship which is still as asymmetrical as that of bomber to bombee.” It's

bad enough that this asymmetry persists in consumer electronics. It's especially discouraging to see leaders in the computer business—Intel, Dell, Microsoft, Apple and even Google—go Hollywood on us.

I heard about the user-side revolution only from bloggers, podcasters and wiki writers, most of whom are glad to take advantage of Linux, open source, free software and free markets. Consumer electronics is gradually being transformed by all these developments, even as its largest brand names still labor, with success, to herd consumers into proprietary silos and walled gardens, now guarded by shiny new DRM systems.

But the revolution will soon become undeniable, even if it isn't televised.

When I first started talking with Andrew, seven years ago, PenguinRadio was a hardware start-up, making radios for playing .mp3 streams. The company still does that, but now most of its revenue comes from advertising in its podcast directory (podcastdirectory.com). Weekly visitors have gone from 150 to 3 million.

A good starting point for measuring the growth of podcasting appears in “DIY Radio With Podcasting”, a piece I wrote for IT Garage (*Linux Journal's* sister site) on September 29, 2004. There, I said, “... now most of my radio listening is to what Adam Curry and others are starting to call podcasts. That last link currently brings up 24 results on Google. A year from now, it will pull up hundreds of thousands, or perhaps even millions.”

I guessed low. According to Wikipedia, “There were 526 hits on September 30, then 2,750 three days later. The number doubled every few days, passing 100,000 by October 18. A year later, Google found more than 100,000,000 hits on the word ‘podcasts’.”

Today (in early January 2006), “blog” brings up 510 million. Perspective: that's 56 million more than the word “consumer”.

“Weblog” brings up 141 million. “Wiki” brings up 240 million.

Compare those to some of the numbers we tracked in UpFront (p. 20), and you'll see how big this transformation is.

It has been my privilege to stand several inches away from ground zero for both the Weblog and podcast explosions. My own blog (doc.weblogs.com) was launched in October 1999, when “blogfather” (of weblogs.com, RSS and too many other developments to name) Dave Winer sat me down and insisted I start blogging. Today my blog is in Technorati's Top 100, out of the 24.7 million blogs (or sources of RSS feeds) tracked by the service. Of the

3.55 million results Google yields for my name, 2.61 also mention "blog". That's compared to 510,000 that also mention "linux".

Technorati was born in November 2003, when David Sifry hacked it up as a research tool to help write the first *Linux Journal* story on blogging. Today Technorati is the #708 Web site on Earth, according to Alexa's traffic rankings. That's ahead of the A9.com search engine, which (like Alexa) is owned by Amazon.com. Technorati is searched several dozen million times a day. (Disclosure: I'm on the Technorati advisory board.)

As for podcasting, I enjoy membership in the Gillmor Gang, a popular weekly podcast that began in 2004. Steve Gillmor's gang takes an hour or so of my time each week. Blogging takes more, but mostly because my efforts have spread to IT Garage, my *SuitWatch* newsletter and *Linux Journal's* own Web site, all of which have RSS feeds and plenty of subscribers. Although my personal blog might run up to a thousand or more words a day, it doesn't take much time because I treat it as a form of public e-mail: a kind of "cc:world". In fact, most of what I write in my blog is in response to e-mails. The rest is in response to subscriptions to keyword searches in Technorati, Pubsub, Google's Blogsearch and other engines that operate in the Live Web that updates constantly (rather than from the relatively Static Web of sites that change slowly and aren't syndicated).

I say all this because I think that many people—even some *Linux Journal* readers—still don't know the extremely high leverage blogging, podcasting and wiki writing can provide. If you write something useful, or provocative, that adds substance to the world—and if you link out to others who serve as sources or also have interesting things to say—the results can be amazing. Search for "Saving the Net" or "Linux is a species", and you'll find stuff I wrote (both for *Linux Journal*) at or near the top of the results.

These results aren't due to "search engine optimization", but rather to the fact that I try to write stuff that's useful, funny, moving, productive or otherwise interesting—knowing that others will want to write about the same things.

All three of this month's topics—blogs, podcasts and wikis—are extraordinarily useful levers on the world. One big reason is that they're personal. Back in the early days of Weblogging, Dave Winer described a blog as "the unedited voice of an individual".

Blogs and podcasts (and even wikis in some cases) can make us much more valuable as employees as well. My roommate at CES was Robert Scoble, a friend of many years who has recently become Microsoft's most well-known blogger, with around 25,000 readers a day. A few days before we got together at CES, Robert raised eyebrows by lambasting his employer for taking down a blogger under pressure from the Chinese government. Speaking frankly to, as well as for, his employer, has been good for both Robert and Microsoft.

Yet, we still tend to see companies as the main instruments of progress, even when the subject is open source.

Feed a cold, starve a fever...
better yet:

Avoid a virus. Use BitDefender!

At least that's what
41 million of our users say.

BitDefender is an award-winning producer of anti-virus software and data security solutions providing protection to satisfy the requirements of today's computing environment.

- **Best reaction to malware as certified by: Andreas Marx (www.av-test.org)**
- **Available for mail and file servers**
- **Answers to critical administrator needs: performance, stability and compatibility**
- **Compliant with FreeBSD 4 & 5 as well as all popular Linux distros**

Now you should feel
quite a BIT better.

 **bitdefender**[®]
SECURE your every bit

Try it for free | www.bitdefender.com

Visit us at Linux World in Boston, April 4-6, 2006 Booth #1234
For a FREE copy of our Linux and FreeBSD Security Solutions



For example, take a *BusinessWeek* article (www.businessweek.com/technology/content/dec2005/tc20051228_262746.htm) from December 28, 2005. The title says “A Watershed for Open Source In 2005”, and the teaser subhead says “The software movement finally gained traction in Corporate America and saw a new influx of VC cash. How will 2006 shape up?” The article lists and describes the “five biggest open-source events of 2005”. The five involve 1) Red Hat, 2) Sun Microsystems, 3) Motorola, 4) Firefox and 5) venture capitalists.

If I had to name five watersheds, I would list developments instead of events. And I would look at what developers and users are doing together over a long period of time, rather than what companies, funders and projects happened to do in the year just past.

Three of my top five developments would be blogs, wikis and podcasts. Not sure what the other two would be. Why? Because I’ve heard Linus and the kernel developers say, “That’s user space. I don’t do user space.” So my natural response is to say “That’s not user space. I only do user space.”

The fun thing about blogs, wikis and podcasts is that users and developers work closely together. In fact, that interaction is essential to progress. For evidence, look at the collaboration around microformats, tagging, structured blogging and OPML. The standards and practices of blogs, podcasts and wikis are all being pushed forward by individuals and developers, working together.

In fact, my favorite explanation of what’s good about the GPL came from Mark Pilgrim, in a blog post titled “Freedom 0” (diveintomark.org/archives/2004/05/14/freedom-0), written shortly after Six Apart came up with a restrictive new license for Movable Type, its formerly (somewhat) open-source blogging software. One excerpt:

Many people misunderstand Free Software and the GNU General Public License. Many people equate the GPL to the boogeyman, because it’s “viral”, and that sounds like a bad thing. Here’s what viral licensing means: GPL software has the restrictions that it has, and that’s it. The GPL is quite restrictive on developers, not at all on end users. (More on that in a minute.) Regardless, GPL software has the restrictions that it has, but it can never become more restrictive. An upgrade can’t take away freedoms that I enjoyed with an older version.

A side effect of this is that if I write a GPL program and then lose interest, and someone else picks it up and continues development, they are forced to release their version under the GPL. A new developer can’t take away freedoms that I enjoyed with the old version either.

I mention this because it’s exactly what happened with WordPress. It started life as b2, which was abandoned. But a year ago, a new community coalesced around a fork of the original b2, and it became WordPress. The new community included some of the original developers, and many new developers. Because the original software was GPL-licensed, WordPress was also GPL-licensed....

...I will never be surprised by the licensing of new versions of WordPress.

Freedom 0 is the freedom to run the program, for any purpose. WordPress gives me that freedom; Movable Type does not. It never really did, but it was “free enough” so we all looked the other way, myself included. But Movable Type 3.0 changes the rules, and prices me right out of the market. I do not have the freedom to run the program for any purpose; I have only the limited set of freedoms that Six Apart chooses to bestow upon me....

WordPress is free software. Its rules will never change. In the event that the WordPress community disbands and development stops, a new community can form around the orphaned code. It’s happened once already. In the extremely unlikely event that every single contributor (including every contributor to the original b2) agrees to relicense the code under a more restrictive license, I can still fork the current GPL-licensed code and start a new community around it. There is always a path forward. There are no dead ends.

Today, a search for “WordPress” brings up 58 million results on Google. A search for “Movable Type” brings up 46.3 million.

I give enormous credit to Mena and Ben Trott, who wrote Movable Type and founded Six Apart, which today hosts millions of MT-based TypePad blogs. The fact that Movable Type was “free enough” helped launch and continues to grow the blogging movement. I also give credit to Google’s Blogger and Userland’s Radio Userland (which runs on Manila, which is written in Frontier, a scripting environment created by Dave Winer, who open sourced it with the GPL in 2004).

Yet the bulk of development work around blogging (including work on syndication, tagging, outlining and other standards and practices) happens outside the corporate context. Matt Mullenweg is better known for his work with WordPress than for whoever his employer happens to be.

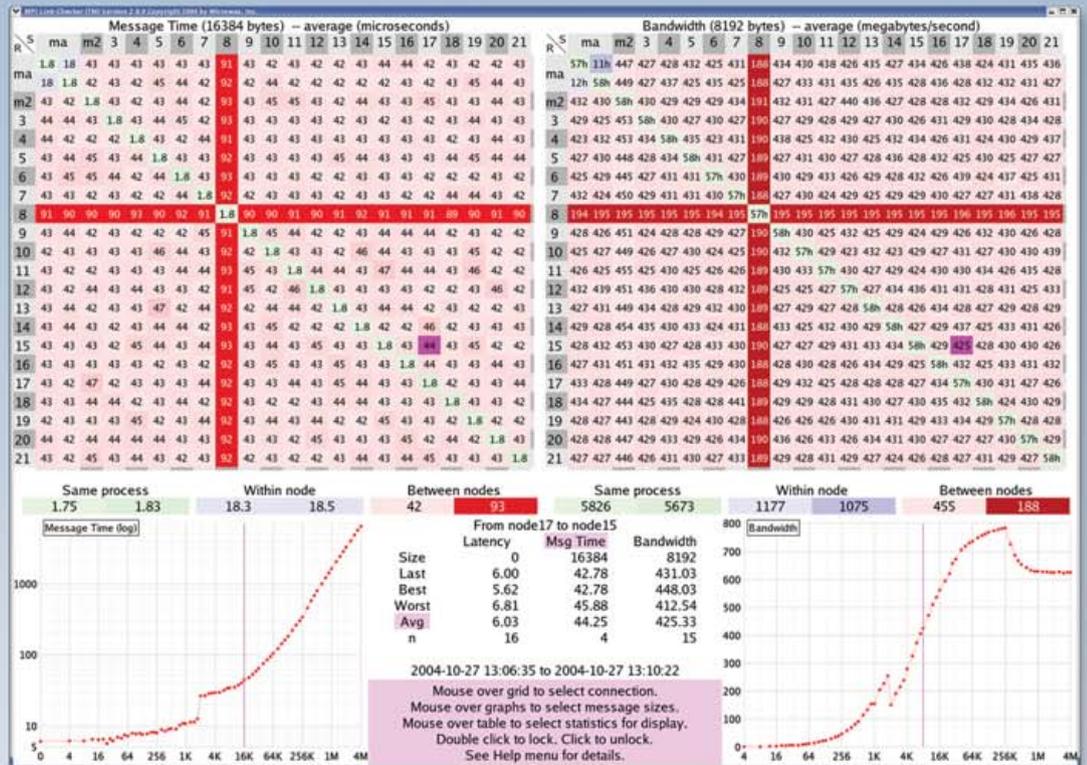
Here’s a revealing fact: the *Wall Street Journal’s* new blogs (blogs.wsj.com) run on WordPress.

There’s a reason this column is called Linux for Suits rather than Linux for Companies. It’s because Linux is about the people who write and use it, not about companies. Note the distinction Linux kernel hackers make between “kernel space” and “user space”. My beat here is the business corner of user space. I’m more interested in what people do with Linux at companies than in what “Linux Companies” are up to.

Blogs, wikis and podcasts grow naturally through the contributions of countless individuals in an environment built and enriched by work on free and open code. Non-free and non-open code can flourish there too. But it’s important to remember where this environment came from. Plenty of credit is due to companies. But far more is due to individuals. It’s what they produce that matters most. Not what the rest of us consume. ■

Doc Searls is Senior Editor of *Linux Journal*.

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 Optron 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 Optron 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 Optrons, 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.



▲ Microway® Quad Optron™ Cluster with 36 Optron 880s, redundant power, 45 hard drives and Myrinet™ in our CoolRak™ cabinet.





LaCie's GlobeTrotter 2.0

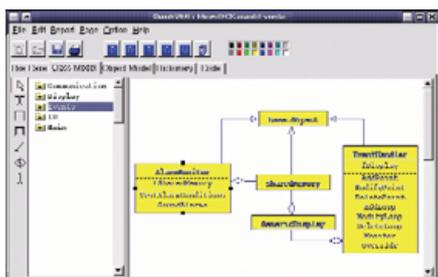
If you need more capabilities than a bootable Flash or CD distribution of Linux offers, LaCie has released a bootable Linux distribution pre-installed onto a 40 or 80GB USB hard drive. Based on its FA Porsche-designed USB hard drive family, the GlobeTrotter 2.0 features a copy of Mandriva Linux Limited Edition 2005 pre-installed on the drive and tuned to launch on just about any system that supports booting from USB devices. The 40GB version retails for \$139.99 US and the 80GB for \$199.99 US, and it includes a one-month membership to Mandriva Online. The drives also include OpenOffice.org.

www.lacie.com

CRYPTOCARD-Server Version 6.3

With government spooks evidently reading all our e-mail, perhaps it's time to invest in a cryptographic solution for your sensitive data. CRYPTOCARD-Server version 6.3 offers an all-in-one access solution combining a smart card or software token with a user-specific pin to validate access to the corporate network. Compatible with SUSE Enterprise Server and Red Hat Enterprise Linux 3, the CRYPTOCARD solution can lock the doors to authorized users trying to sneak in via PPP, SSH, Samba and VPN solutions, as well as the Linux display managers. A five-user package with the server and five tokens is available for \$499 US.

www.cryptocard.com



Excel Software's QuickUML 2.0

Does your job require you to draw pretty pictures with boxes and lines before you get to write any code? If UML is party of your workplace reality, Excel Software has begun shipping QuickUML 2.0 for Linux, Windows and Mac. According to Excel, QuickUML is designed to be, well, quick.

They claim that no formal training is needed to begin designing objects using the tool, and that QuickUML can generate code stubs for C++, Java or Delphi. Single-user licenses start at \$295 US.

www.excelsoftware.com

IceWarp Software Ltd.'s Merak Mail Server

You've got mail. Lots and lots of mail. Maybe you need the Merak mail server from IceWarp Software Ltd. Already a widely deployed server under Windows, the new Linux edition of the software is specifically tuned to the OS. Certified for use with Red Hat Enterprise Linux, it also has been tested on Red Hat Fedora Core and Gentoo. Success also has been reported on SUSE, Mandrake, Ubuntu and Debian distributions. The package includes antispam, antivirus, groupware and instant-messaging capabilities, and it is designed for loads of up to 20,000 messages per minute and up to 100,000 active accounts. Pricing runs from \$210 US for a single domain version to \$5,700 US for a full-blown package with all the features.

www.icewarp.com

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

Why is LPI the World's #1 IT certification?

Stable.

All Linux Professional Institute certification programs are created using extensive community input, combined with rigorous psychometric scrutiny and professional delivery. We test the whole continuum of important Linux skills - we don't just focus on small, subjective tasks.

Innovative.

LPI's program follows the LSB specification, so people who pass our tests can work on all major distributions. Because of its strong grassroots base and corporate support both inside and outside the world of open source, LPI goes beyond "vendor-neutral" to truly address community needs.

Growing.

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



**Linux
Professional
Institute**

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



Why I Don't Like [edit: Get] Wikis

Without editorial control
[edit: tyrannical egotistic
editors in chief] wikis
are at best a good idea
done bad.

Dave Taylor

I've been called a curmudgeon before [edit: because you are] and there are some technologies [edit: like television] that I don't really see as astonishingly useful evolutionary steps in the world of information and technology, but even with that disclaimer, I have to say that I'm completely unimpressed with wikis and really don't understand why so many other people love them so [edit: maybe because they're just smarter than you are, jerk!].

Intellectually, the idea of collaborative editing and maintenance of text documents is quite appealing, but the pragmatic reality of having essentially zero editorial control over content is problematic at best and dangerous at worst. Would you trust a medical encyclopedia built around wiki technology?

~~But let's start by defining a wiki, shall we?~~ [edit: These sort of rhetorical questions are just trite author tricks and should be axed.]

A wiki, which gets its name from the Hawaiian phrase "wiki wiki", or "quick, quick", is a simple software application that allows content to be separated from its presentation to make it trivially easy to have visitors modify and change any content that they see [edit: on a wiki-based Web site]. ~~More sophisticated~~ [add: Just about all] wiki packages have a sophisticated, if arcane [edit: it's not arcane, you're just stupid] [edit: come on, when =a= and ==a== produce different formatting, it's pretty arcane] [edit: screw you, l0s3r] markup language.

If this article were hosted on a wiki [edit: too bad it's not. Then we could fix all the inane commentary herein], you could decide that you don't like my definition of the word "wiki", or even the article title, click on an edit button and change things to your heart's content. There's a change-tracking mechanism built into all wiki systems (and it should be no surprise that's a critical element [edit: if only to get rid of stupid edits]), but you can imagine that when pages can be edited and modified five, ten or even 20 or more times daily, it can ~~lead to a painful editorial management task~~ [add: be almost impossible to retain any sort of quality control over the content]. [edit: The point of a wiki is that there isn't any editorial control, though. This entire premise is false.]

Now, let's say that I wanted to write about the infamous Skull & Bones Society and its intersection with the Illuminati, Opus Dei and the Bush family [edit: and your mama, too]. You can easily imagine that my take on this vast conspiracy might well be dramatically different from your take, and sure enough, there are certain types of content that ~~really~~ suffer the worst in wikis, as the on-again, off-again article on JFK's assassination on Wikipedia demonstrates. It seems that a ~~crackpot~~ [add: guy who didn't buy the government coverup] decided that there was a conspiracy involved in Kennedy's assassination and added that to the page. But others felt otherwise and purged

the Wikipedia entry of his content. And he added it back. And they deleted it. To the point where it's now impossible to know whether the page reflects the commonly held facts of the situation or some crank theory. [edit: Truth is subjective.]

Even with smaller groups, I've tried having a wiki for a team of about a dozen people, and the necessity of using the arcane wiki coding schemes and confusion of tracking edits rapidly diminished anyone's enthusiasm for the new technology and the project quickly ran out of steam. [edit: You were probably all just too st00pd to use a wiki!] Document tracking in Microsoft Word is far, far easier, and it's not that hard to e-mail files around, even in this day and age [edit: and horses and buggies? Is that your speed too?].

I suppose wikis have their place and

certainly there are fans who find them a useful Web-based document "evolution" petri dish, if you will. [edit: Sheesh, can we PLEASE purge this guy of his clichés? This is a terrible article!] [edit: Yeah, and what have you written and published lately, chump?] [edit: Where is that relevant, l0s3r?] [edit: Can't you just GO AWAY and leave this page alone?] [edit: I will when it's accurate] [edit: According to who, you?] [edit: Hey, I can edit this more than you can. Wanna test me?]

Overall, though, the only time I have seen wikis work is when not everyone who wanders onto the site can edit the content, but if there's editorial control, it seems to be counter to the basic premise of wikis, that they're a tool for leveraging the collaborative editorial efforts of the public.

That's why I believe that as technologies go, wikis are going to end up in the good idea, bad implementation, or, perhaps, good concept, bad fit with reality graveyard. [edit: That's okay, you'll be there too, Taylor, and this article shows exactly why.]■

Dave Taylor has been involved with UNIX and Internet technologies since 1980 and has picked some winners in the technology sweepstakes (even in 1980 it was clear that e-mail was the killer app for networks), but backed some clunkers too. You can pick up the debate on his business blog *The Intuitive Life*, at www.intuitive.com/blog.

Ultra Dense, Powerful, Reliable... Datacenter Management Simplified! 15" Deep, 2-Xeon/Opteron or P4 (w/RAID) options



Customized Solutions for... Linux, BSD, W2K

High Performance Networking Solutions

- Data Center Management
- Application Clustering
- Network and Storage Engines

Rackmount Server Products

- 1U Starting at \$499: C3-1GHz, LAN, 256MB, 20GB IDE
- 2U with 16 Blades, Fast Deployment & more...



iron
SYSTEMS™

Iron Systems, Inc.

540 Dado Street, San Jose, CA 95131
www.ironsystems.com

CALL: 1-800-921-IRON

Using Wikis and Blogs to Ease Administration

This tutorial on TWiki and WordPress shows how wikis and blogs can be useful for system administration and documentation. **TI LEGGETT**

System administration can be like sailing a ship. You must keep your engines running smoothly, keep your crew and the harbors notified and up to date and also maintain your Captain's log. You must keep your eye on the horizon for what is coming next. Two technologies have emerged over the past few years that could help keep you on course, wikis and blogs.

Maintaining Good Documentation

I find that one of the most difficult aspects of system administration is keeping documentation accurate and up to date. Documenting how you fixed a pesky problem today will help you remember how to fix it months later when it occurs again. If you ever have worked with others, you realize how critical good documentation is. Even if you are the only system administrator, you still will reap the benefits of good documentation, even more so if another sysadmin is ever brought on board.

Some goals of a good documentation system should be:

- Make it easy for you and your coworkers to find relevant information.
- Make it easy for new employees to come up to speed quickly.
- Make it easy to create, edit and retire documentation.
- Keep revisions of changes and who made them.
- Limit who sees or edits the documentation with an authentication system.

Unfortunately, keeping your documentation up to date can be a full-time job in itself. Documenting, though not a very glamorous task, certainly will pay off in the long run.

Why a Wiki?

This is where a wiki comes in. From Wikipedia: "a wiki is a type of Web site that allows users to add and edit content and is especially suited for constructive collaborative authoring."

What this means is a wiki allows you to keep and edit your documentation in a central location. You can access and edit that documentation regardless of the platform you are using. All you need is a Web browser. Some wikis have the ability to keep track of each revision of a changed document, so you can revert to a previous version if some errant changes are made to a document. The only obstacle a new user must overcome is learning the particular markup language of your wiki, and sometimes even this is not completely necessary.

One of a wiki's features is also one of its drawbacks. Wikis are pretty free flowing, and although this allows you to concentrate on getting the documentation written quickly, it can make organization of your wiki rapidly spiral out of control. Thought needs to be put into how the wiki is organized, so that topics do not get stranded or lost. I have found that making the front page a table of contents of all the topics is very handy. However you decide to organize your wiki, make sure it is well understood by everyone else. In fact, a good first document might be the policy describing the organization of the wiki!

TWiki

There are several open-source wikis available, such as MediaWiki [see Reuven M. Lerner's article on page 62 for more information on MediaWiki] and MoinMoin, each with its own philosophy on markup and layout, but here we concentrate on TWiki. Some of TWiki's benefits are:

- A notion of webs that allows the wiki administrator to segregate areas of collaboration into their own areas, each with its own set of authorization rules and topics.
- A modular plugin and skin system that allows you to customize easily.
- A well-established base of users and developers.
- Revision control based on RCS.
- It is Perl-based and mod_perl or FastCGI can be used.
- Authentication is handled outside the wiki by mechanisms such as Apache httpasswd.

The most current stable release at this time is Cairo, or TWiki20040904. It was released, as the name suggests, on September 4, 2004, and it has been proven to be very stable. However, it does lack some of the features of the current beta release, Dakar, that I find to be very useful. The Dakar release we use here is TWikiRelease2005x12x17x7873beta.

Installing TWiki is relatively easy, but still needs work. I hope, as the beta progresses, we will see improvements in ease of installation and upgrading along with clearer documentation.

First, you must create the directory where you want to install TWiki, say /var/www/wiki. Next, untar the TWiki distribution in that directory. Then you must make sure that the user with rights to run CGI scripts (usually apache or www-data), owns all of the files and is able to write to all files:

```
# install -d -o apache /var/www/wiki
# cd /var/www/wiki
# tar xzf /path/to/TWikiRelease2005x12x17x7873beta.tgz
# cp bin/LocalLib.cfg.txt bin/LocalLib.cfg
# vi bin/LocalLib.cfg lib/LocalSite.cfg
```

```
# chown -R apache *
# chmod -R u+w *
```

Now copy bin/LocalLib.cfg.txt to bin/LocalLib.cfg, and edit it. You need to edit the \$twikiLibPath variable to point to the absolute path of your TWiki lib directory, /var/www/wiki/lib in our case. You also must create lib/LocalSite.cfg to reflect your specific site information. Here is a sample of what might go into LocalSite.cfg:

```
# This is LocalSite.cfg. It contains all the setups for your local
# TWiki site.
$cfg{DefaultUrlHost} = "http://www.example.com";
$cfg{ScriptUrlPath} = "/wiki/bin";

$cfg{PubUrlPath} = "/wiki/pub";
$cfg{DataDir} = "/var/www/wiki/data";
$cfg{PubDir} = "/var/www/wiki/pub";
$cfg{TemplateDir} = "/var/www/wiki/templates";
$TWiki::cfg{LocalesDir} = '/var/www/wiki/locale';
```

Here is a sample section for your Apache configuration file that allows this wiki to run:

```
ScriptAlias /wiki/bin/ "/var/www/wiki/bin/"
Alias /wiki "/var/www/localhost/wiki"
<Directory "/var/www/wiki/bin">
    Options +ExecCGI -Indexes
    SetHandler cgi-script
    AllowOverride All
    Allow from all
</Directory>
<Directory "/var/www/wiki/pub">
    Options FollowSymLinks +Includes
    AllowOverride None
    Allow from all
</Directory>
<Directory "/var/www/wiki/data">
    deny from all
</Directory>
<Directory "/var/www/wiki/lib">
    deny from all
</Directory>
<Directory "/var/www/wiki/templates">
    deny from all
</Directory>
```

TWiki comes with a configure script that you run to set up TWiki. This script is used not only on initial install but also when you want to enable plugins later. At this point, you are ready to configure TWiki, so point your browser to your TWiki configure script, <http://www.example.com/wiki/bin/configure>. You might be particularly interested in the Security section, but we will visit this shortly. Until you have registered your first user, you should leave all settings as they are. If the configure script gives any warnings or errors, you should fix those first and re-run the script. Once you click Next, you are prompted to enter a password. This password is used whenever the configure script is run in the future to help ensure no improper access.

Once you have completed the configuration successfully, it is time to enter the wiki. Point your browser to <http://www.example.com/wiki/bin/view>, and you are presented with the Main web. In the middle of the page is a link for registration. Register yourself as a user. Be sure to provide a

valid e-mail address as the software uses it to validate your account. Once you have verified your user account, you need to add yourself to the TWikiAdminGroup. Return to the Main web and click on the Groups link at the left, and then choose the TWikiAdminGroup. Edit this page, and change the GROUP variable to include your new user name:

```
Set GROUP = %MAINWEB%.TiLeggett
Set ALLOWTOPICCHANGE = %MAINWEB%.TWikiAdminGroup
```

The three blank spaces at the beginning of each of those lines are critical. These two lines add your user to the TWikiAdminGroup and allow only members of the TWikiAdminGroup to modify the group. We are now ready to enable authentication for our wiki, so go back to <http://www.example.com/wiki/bin/configure>. Several options provided under the Security section are useful. You should make sure the options {UseClientSessions} and {Sessions}{UseIPMatching} are enabled. Also set the {LoginManager} option to TWiki::Client::TemplateLogin and {PasswordManager} to TWiki::Users::HtPasswdUser. If your server supports it, you should set {HtPasswd}{Encoding} to sha1. Save your changes and return to the wiki. If you are not logged in automatically, there is a link at the top left of the page that allows you to do so.

Now that you have authentication working, you may want to tighten down your wiki so that unauthorized people do not turn your documentation repository into an illicit data repository. TWiki has a pretty sophisticated authorization system that is tiered from the site-wide preferences all the way down to a specific topic. Before locking down the Main web, a few more tasks need to be done. Once only certain users can change the Main web, registering new users will fail. That is because part of the user registration process involves creating a topic for that user under the Main web. Dakar has a user, TWikiRegistrationAgent, that is used to do this. From the Main web, use the Jump box at the top left to jump to the WebPreferences topic. Edit the topic to include the following four lines and save your changes:

```
Set ALLOWTOPICRENAME = %MAINWEB%.TWikiAdminGroup
Set ALLOWTOPICCHANGE = %MAINWEB%.TWikiAdminGroup
Set ALLOWWEBRENAME = %MAINWEB%.TWikiAdminGroup
Set ALLOWWEBCHANGE = %MAINWEB%.TWikiAdminGroup,
  ➤%MAINWEB%.TWikiRegistrationAgent
```

This allows only members of the TWikiAdminGroup to make changes or rename the Main web or update the Main web's preferences. It also allows the TWikiRegistrationAgent user to create new users' home topics when new users register. I have included a patch that you must apply to lib/TWiki/UI/Register.pm as well. The patch follows, but you can also download the patch from the [LJ](#) FTP site (see the on-line Resources):

```
--- lib/TWiki/UI/Register.pm.orig      2006-01-04 01:34:48.968947681 -0600
+++ lib/TWiki/UI/Register.pm          2006-01-04 01:35:48.999652157 -0600
@@ -828,11 +828,12 @@

     my $userName = $data->{remoteUser} || $data->{WikiName};
     my $user = $session->{users}->findUser( $userName );
+   my $agent = $session->{users}->findUser( $twikiRegistrationAgent );
     $text = $session->expandVariablesOnTopicCreation( $text, $user );

     $meta->put( 'TOPICPARENT', { 'name' => $TWiki::cfg{UsersTopicName} } );

-   $session->{store}->saveTopic($user, $data->{webName},
+   $session->{store}->saveTopic($agent, $data->{webName},
                                   $data->{WikiName}, $text, $meta );
```

```
    return $log;
}
```

Otherwise, new users' home directories will fail to be created and new user registration will fail. Once you have verified that the Main web is locked down, you should do the same for the TWiki and Sandbox webs.

When you are done configuring TWiki, you should secure the files' permissions:

```
# find /var/www/wiki/ -type d -exec chmod 0755 {} ';'
# find /var/www/wiki/ -type f -exec chmod 0400 {} ';'
# find /var/www/wiki/pub/ -type f -exec chmod 0600 {} ';'
# find /var/www/wiki/data/ -type f -exec chmod 0600 {} ';'
# find /var/www/wiki/lib/LocalSite.cfg -exec chmod 0600 {} ';'
# find /var/www/wiki/bin/ -type f -exec chmod 0700 {} ';'
# chown -R apache /var/www/wiki/*
```

As I mentioned before, TWiki has a plugin system that you can use. Many plugins are available from the TWiki Web site. Be sure the plugins you choose have been updated for Dakar before you use them.

Keeping Your Users in the Know

One important aspect of system administration that is sometimes overlooked is keeping users informed. Most users like to know when there is new functionality available or when resources are down or not available. Not only does it make users happier to be kept informed, but it also can make your life easier as well. The last thing you want to do when the central file server is down is reply to users' questions about why they cannot get to their files. If you have trained your users to look at a central location for status of the infrastructure first, all you have to do after notification of a problem is post to this central place that there is a problem. Mailing lists also are good for this, but what if the mail server is down? Some people, for instance your boss or VP of the company, might like to know what the status is of things as they happen. These updates might not be suitable to send out to everyone daily via e-mail. You could create yet another mailing list for these notifications, but you also might consider a blog.

If you are not familiar with a blog, let us refer back to Wikipedia: "a blog is a Web site in which journal entries are posted on a regular basis and displayed in reverse chronological order."

The notion of a blog has been around for centuries in the form of diaries, but blogs recently have had an explosion on the Internet. Many times a blog is started as someone's personal journal or as a way to report news, but blogs can be extremely useful for the sysadmin.

Blogs can help a sysadmin give users an up-to-the-minute status of what they are doing and what the state of the infrastructure is. If you faithfully update your blog, you easily can look back on what you have accomplished so you can make your case for that raise you have been hoping for. It also will help you keep track of what your coworkers are doing. And, with many blog software packages providing RSS feeds, users can subscribe to the blog and be notified when there are new posts.

WordPress

There are a lot of blog software packages out there today, but here we cover WordPress. WordPress is fast and has a nice plugin and skin interface to allow you to customize it to your heart's content. The only requirements for running WordPress are Apache, MySQL and PHP. I don't go into how to install WordPress, because the on-line documentation is very clear and easy to follow. Instead, I start where the installation leaves off and introduce some useful plugins. I suggest starting with WordPress v1.5.2 even though v2.0 is currently out. There have been some problems with the initial 2.0

release that warrant waiting for v2.0.1. Also, many of the plugins have not had a chance to update to the new system.

The first thing you should do after installing WordPress is log in as the admin user. Once logged in, you are presented with the Dashboard. At the top of the page is a menu of options named Write, Manage, Links and so on. You should first create an account for yourself by clicking on the Users option. Once that has loaded, two tabs labeled Your Profile and Authors & Users are available under the main menu. Click on Authors & Users, and scroll down to the Add New User section and fill in the text fields. Once your user has been added, it appears in the Registered Users section above. There are several columns of data, and one is Promote, which you should click on. Promoting a user makes that user an author and also allows that user to have more privileges based on its level. Once your user has been promoted, it will have a level of one. There are plus and minus signs on either side of the level to use to increase your user's level. Increase it to nine, which is the highest level a non-admin user can be. Should you ever need to delete users that have been promoted to authors, all you need to do is decrease their level below one and then delete them. I have included a link to a more in-depth description of the privileges of each user level in the on-line Resources.

There are a few other options you might consider changing. In General Options, there are check boxes to allow anyone to register to become a blog user and to require users to be logged in to add comments. You may or may not want these options enabled, depending on your security concerns and the openness of your blog. At our site, users cannot register themselves, though anyone can post comments without being logged in. You should explore all the menus and all their options to tweak them for your site.

WordPress Plugins

WordPress has a very modular plugin system, and a lot of people have written many plugins. WordPress also has a notion of categories. Categories can have many uses, but one might be to create mini-blogs for different communities of users or to group posts about a specific aspect of the infrastructure. But, you might not want all users to be able to see every category. The Userextra plugin, in conjunction with the Usermeta plugin, allows you to control exactly this sort of thing. Once you have followed these plugins' installation instructions, two more menus are available under Options and one more under Manage that allow you to refine access.

Another plugin you may find useful is the HTTP Authentication plugin. This plugin lets you use an external authentication mechanism, such as Apache's BasicAuth, as a means to authenticate to WordPress. This is great if you already have an LDAP directory or Kerberos realm that you use for authentication and you have mod_auth_ldap or mod_auth_kerb up and running.

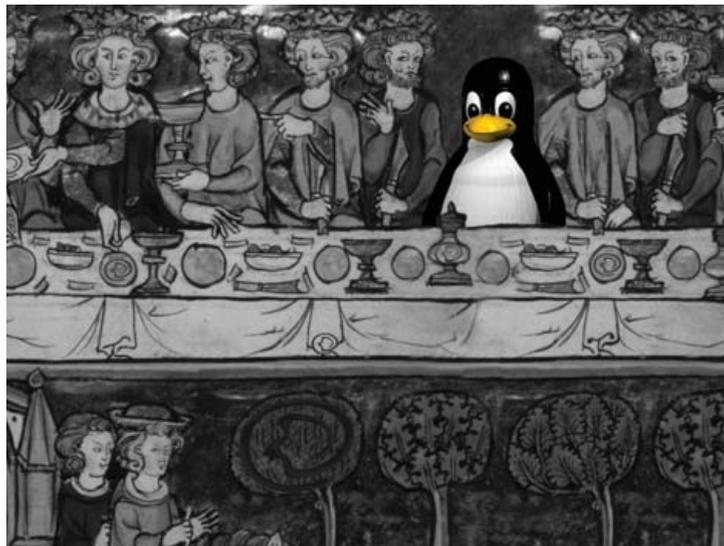
Many more plugins are available for WordPress from the WordPress Codex and the WordPress Plugin DB. If you feel some functionality is missing, there are plenty of examples and documentation available from the WordPress Web site, and these plugin repositories can help you write your own plugin.

Wrapping Up

I hope that after this whirlwind tour of wikis and blogs you have come to see how they can be beneficial to help your shop run a smoother ship and provide your users with all the information they might want. Just as there are many different sails to keep your ship sailing, there are many different wiki and blog software packages out there. The right package for you is the one that keeps your users happy and you productive.

Resources for this article: www.linuxjournal.com/article/8832. ■

Ti Leggett (ti@daleggetts.com) is a full-time system administrator. When he's not working, he might be found playing his Gibson B-25 or doing some home improvements or wood working.



The holy GRAIL

Easy, high-performance clustering. For years, many searched, but none could find it. Some said it didn't exist. But not the Penguin.

Penguin Computing® made easy, high-performance clustering a quest. Now you can find dramatically simplified Linux clusters, deployed and managed from a single point, driven by Scyld's commercially supported, industry-leading Linux clustering software and the high efficiency AMD Dual Core Opteron™ for improved productivity. For the turnkey clusters, needed to run even your most important applications, come to Penguin Computing.

Powerful, easy clustering. It's the once and future thing. **Love what you do.** 

www.penguincomputing.com

Join us at Linux World at the Boston Convention & Exhibition Center, Booth # 1324



Penguin Computing is a registered trademark of Penguin Computing, Inc. Scyld Software and the Scyld Block Logo are registered trademarks of Scyld Software, Inc. Linux is a registered trademark of Linus Torvalds. AMD, AMD Opteron, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. Other names are for informational purposes only and may be trademarks of their respective owners.
Copyright © 2000-2006 Penguin Computing, Inc. All rights reserved.



Installing and Customizing MediaWiki

All you need to know to install and use the powerful MediaWiki system. Reuven M. Lerner

I was first introduced to the idea of a wiki many years ago. A colleague at work told me about a Web site that invited readers to become writers and editors. Over time, he said, a wiki would grow to include a great deal of knowledge, drawing from the collective experience, without the need for any centrally coordinated assignment or editing. My gut reaction was one of great ambivalence. On the one hand, a community of well-meaning and well-disciplined users could indeed use a wiki in many useful ways. But at the same time, it seemed like an invitation for chaos. And, admittedly, chaos has sometimes been the result.

Over the years, wikis have become increasingly common. Dozens of different software packages are now available, in nearly every programming language and for nearly every Web/database framework. The question is no longer whether a wiki is available, or how easily one can be installed, but rather whether a wiki is an appropriate tool for your needs and which of the existing packages is the best fit for your organization.

And although the jury is still out, at least in part, regarding the long-term viability of public wikis, the amazing Wikipedia Project has demonstrated that it is possible to have surprisingly good results. A research study released several weeks before I wrote this article reported that when it comes to science-related articles, Wikipedia is only slightly less accurate than Encyclopædia Britannica. Given that Britannica is written and edited by paid experts, while Wikipedia relies on volunteers (who may, of course, still be experts), it is clear that wikis do not necessarily lead to the chaos that I feared.

That said, I have a strong feeling Wikipedia's reputation was pre-

served not because of the public's inherent love for wikis, but because of a dedicated team of Wikipedians who operate behind the scenes to ensure that the content is viable.

As I wrote above, there are many options for creating a wiki on your own Web server. One of the best-known packages is the same one that powers Wikipedia, known as MediaWiki. In this article, we discuss how you can install MediaWiki on your own server, how you can create and edit static content and how you can even create dynamic special pages.

Installing MediaWiki

Although Wikipedia contains an enormous amount of text, the MediaWiki software on which it runs is surprisingly small, straightforward to understand and easy to install. The software itself is written in PHP, and it relies on a MySQL database for content and indexing. That is, none of the pages on a MediaWiki site exist as files on disk; instead, they are created on the fly by the PHP programs, bringing together a number of elements for each individual visitor.

To install the MediaWiki software, you thus need a server running both PHP and MySQL. PHP works on a number of platforms, but we will assume, for the purposes of this article, that you are using the Apache server with relatively recent versions of PHP (at least 4.1.2, although 4.3 is preferred) and MySQL (at least 3.2.x, but 4.0.x is preferred). Most modern Linux distributions either include PHP and MySQL or make them available from on-line repositories without too much trouble. You need administrative privileges for both Apache and MySQL in order to install MediaWiki. If

you don't run your own server, you might need to ask the system administrator for some help in modifying the appropriate configuration files.

Although you can download the MediaWiki software directly from SourceForge, you should look for it on the main MediaWiki site, www.mediawiki.org. Once on that page, click on the versions and download link (on the right side of the page), and then on the appropriate link for the current stable release. Download an older or experimental release only if you understand the problems that might be associated with it.

The software comes as a .tar.gz file containing the version number. For example, the file that I downloaded is called mediawiki-1.5.3.tar.gz, indicating that I downloaded version 1.5.3, current as of December 21, 2005. Keep that file in a temporary location (I generally prefer to put such files in a directory named /downloads); we will return to it after taking care of our Apache configuration.

If your site will run only MediaWiki software, there is probably no reason for you to modify your Apache configuration file, traditionally named httpd.conf. In such cases, you can unpack the .tar.gz file in your site's htdocs directory, as described in the instructions on the MediaWiki site. If your distribution uses another directory as its document root, you'll obviously have to adapt to your distribution's preferences.

Virtual Servers

If you want to put the MediaWiki installation in a virtual server on an existing Apache installation, you need to modify the Apache configuration file to create a new virtual host on your system. Modern versions of HTTP require that a Web browser request a document from a specific hostname. The virtual host feature in Apache allows one Web server to serve documents for many different hostnames, even when all of the names resolve to a single IP address. If you just want to try MediaWiki without affecting the rest of your site, or if the majority of your site runs under a framework (such as Zope or Ruby on Rails) that is largely incompatible with MediaWiki, a virtual host might well be a wise move. I often use such virtual hosts to experiment with new software and configurations without potentially endangering existing, stable sites.

Apache's VirtualHost sections allow you to configure one or more such virtual servers, each with its own configuration. In configuring MediaWiki for my system, I needed the following VirtualHost section in httpd.conf:

```
<VirtualHost 69.55.225.93>

ServerName wiki.lerner.co.il
    ServerAlias mediawiki.lerner.co.il wikipedia.lerner.co.il
ServerAdmin reuven@lerner.co.il

DocumentRoot /usr/local/apache/v-sites/wiki.lerner.co.il/www/
DirectoryIndex index.php

CustomLog
/usr/local/apache/v-sites/wiki.lerner.co.il/logs/access-log combined
CustomLog
/usr/local/apache/v-sites/wiki.lerner.co.il/logs/referer-log referer
ErrorLog /usr/local/apache/v-sites/wiki.lerner.co.il/logs/error-log

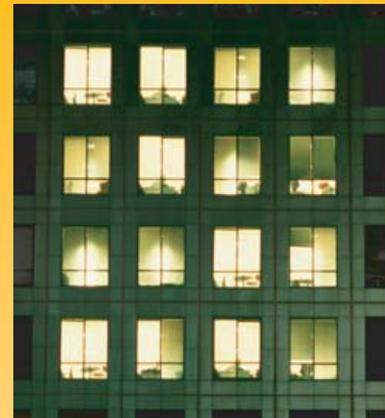
</VirtualHost>
```

The first two lines of this VirtualHost section define the name and aliases for this new virtual server. By adding aliases, I ensure that even if people enter a different name, they will get to the right place. (Because I have asked my hosting company to alias all hostnames in the lerner.co.il domain to my server's IP address, I can add and remove new virtual hosts whenever I want, without having to modify the DNS configuration.) I then set ServerAdmin to be my own e-mail address, ensuring that error messages will direct people to me.

Next, I indicate that the root directory for the wiki—that is, where

SCYLD WORKFORCE

Pronunciation: *skild* (That's a hard "sc" as in "scalability," not a sibilant "sc" as in "sci-fi")
werkfors (sounds pretty much like it looks)
Function: *idiomatic expression*
Etymology: *Scyld*, from Middle English *skilled*, to be exceptionally talented, trained, or abled
Workforce, the people who make the wheels turn and keep the lights on
Usage: See the difference software can make to your workforce. Download "Breaking New Ground: The Evolution of Linux Clustering" at www.scyld.com/hpc.



1: employees who use Scyld Beowulf[®], the Linux clustering software that does it all **2:** engineers, researchers and sysadmins alike, who need powerful yet elegant solutions **3:** highly talented people who focus on managing their jobs, not their clusters **4:** how's this for some turn-key, personnel pleasing features **a:** commercial-grade solution <as in the end of do-it-yourself Linux clustering> **b:** single point of management <as in wickedly simple and highly scalable **5:** all those who don't want to change the way they work, just the results they're used to getting

synonyms: elegance, simplicity, power

antonyms: labor intensive, SMP, Unix, Windows

Join us at Linux World at the Boston Convention & Exhibition Center, Booth # 306

Highly
SCYLD
www.scyld.com

Apache should look for the PHP programs that define it—is under `/usr/local/apache/v-sites`, where I put all of the virtual servers on my machine. Each site then gets a subdirectory according to its name, with two directories (`www` and `logs`) under that. So in the above configuration, I have defined the `DocumentRoot` to be under `v-sites/wiki.lerner.co.il/www`, and I have put the access, referer and error logs under `v-sites/wiki.lerner.co.il/logs`.

Once I have added the above configuration to `httpd.conf`, I restart Apache. At this point, I'll get an error if I go to `http://wiki.lerner.co.il`, as I haven't yet installed the software.

To install the MediaWiki software, I download the latest `.tar.gz` file, and open it in the `v-sites/wiki.lerner.co.il`:

```
cd /usr/local/apache/v-sites/wiki.lerner.co.il/  
tar -zxvf /downloads/mediawiki-1.5.3.tar.gz  
mv mediawiki-1.5.3 www
```

After doing that, I pointed my Web browser to `http://wiki.lerner.co.il`, and got a message saying that I still needed to configure the software. But the permissions weren't set quite right, meaning that MediaWiki was unable to write to its configuration directory. Rather than aborting with a hard-to-understand error message, MediaWiki kindly told me that there was a problem with the filesystem permissions, and it even reminded me how to use `chmod` to fix the problem.

After having changed the permissions, MediaWiki then showed me the initial configuration screen. To be honest, I didn't have to do very much; I entered the name of my wiki (ReuvenWiki), and I kept the defaults for nearly everything else. I entered a password for the WikiSysop user within the wiki system itself. And, I entered a password for the database user and also entered the password for the MySQL administrative user.

Remember, you're keeping track of three types of user names here, which can be a bit tricky. Your Linux system has one set of users and permissions, MySQL has its own set of users and permissions, and the MediaWiki software also has its own set of users and permissions. Although this might seem like overkill, it provides a great deal of flexibility, allowing you to run MediaWiki in a safe, secure environment.

Once I filled out all of the above information, I clicked on the `Install!` button. Unlike many Web applications, which produce output only when they have finished with their work, the MediaWiki installer kindly provides updates as the installer is running, giving you a sense of what it is doing and how much longer you might have to wait. In the end, you should (hopefully) see a note indicating that the installation was successful, and that you should move `config/LocalSettings.php` into the parent directory (the `www` directory):

```
mv -v config/LocalSettings.php .
```

Having done that, I can re-enter the URL (`http://wiki.lerner.co.il`) into my Web browser. And, sure enough, I'm greeted by MediaWiki's initial page, which tells me that the software has been installed successfully and points me to two pages, one for customizing the interface and another describing usage and configuration.

Basic Navigation and Editing

If you have ever used Wikipedia, navigating through a MediaWiki site will be extremely familiar to you. If not, the basic rules are as follows: clicking on a link within a page brings you to another page within the MediaWiki site. The exception is when the link is followed by an arrow icon, which indicates that the link will open a site on another server. All pages on the site are not only viewable by the general public, but they also are editable. If you want to change the contents of a page, click on the `Edit` link (at the top of each MediaWiki screen), modify the contents as you see fit, and then click on the `Save` button when you are done.

MediaWiki tries to make it as easy as possible for users to participate

without having to understand the technical side of things. It's thus unnecessary for end users to know HTML; they instead can click on a variety of buttons (and use a simplified formatting system) to format postings they make to the system. For example, entering `[[Foo]]` creates a link to the page `Foo`, regardless of whether it already exists. And, entering a link of the form `[[Foo | This leads to Foo]]` inserts a link to the `Foo` page, but with the text "This leads to Foo".

More important than MediaWiki's simple markup language is the fact that every change made to the system is kept in an easy-to-use version-control system. By clicking on the `History` link on a given page, even a non-technical user can view previous versions of a page, see who has made a given change or compare two versions of a particular document. In addition, users can set up watch lists, such that the user receives an e-mail message every time a page is modified. This allows users with particular interests and expertise to keep up with the latest updates to a given page and to monitor the content for potential damage.

By default, the left side of every page contains a short menu of commonly accessed links, including the main page (home page) of the wiki, a list of recent changes, a random existing page and the documentation. There is also a search box on the left side. Entering a text string and clicking on the `Search` button brings up a list of documents whose titles and/or contents match the target string. Because MediaWiki contents are all in the MySQL database, and because modern versions of MySQL have built-in text-indexing routines, such searches are relatively speedy.

Each page in a MediaWiki site has a URL that begins with the overall site URL (in my case, that would be `http://wiki.lerner.co.il`), followed by `index.php/`, and then the title of the page. Page titles may contain spaces, in which case the URL replaces each space with an underscore character. The page named `Reuven Lerner` on my server would thus have a complete URL of `http://wiki.lerner.co.il/index.php/Reuven_Lerner`. Somewhat frustratingly, MediaWiki URLs are case-sensitive, which means that entering an incorrectly capitalized word can result in the creation of a new page. Remember, a wiki never produces a true "not found" error, because you (the visitor) might well be interested in starting that very page.

There is no hierarchy to pages in a MediaWiki site. There is a single namespace, which can potentially lead to confusion. To reduce ambiguity, pages may contain additional information in parentheses (which are then part of the URL). Or, users can create a disambiguation page, which contains links to a number of pages with similar or identical names.

If you want to allow users to upload images, you need to modify the `LocalSettings.php` file that should now exist in your server's `DocumentRoot` directory. In the version I installed, line 73 of `LocalSettings.php` (just beneath a relevant comment in the PHP code) contains a commented-out definition of `$wgEnableUploads`, setting it to `true`. By uncommenting this line, you allow users to upload images to the server. These uploads will be, by default, in the `images` subdirectory (`www/images`), whose permissions must be writable by the Apache process.

Once you have turned on the uploading of images, logged-in users will be able to upload them with a two-step process. First, they click on the `insert image` button when editing a wiki page, or manually insert tags of the form `[[Image:Example.gif]]`. (Images all have the `Image:` prefix in MediaWiki, but there is a single, flat namespace for images, as is the case for textual content pages.) Once users have finished editing the page, they can then click on the missing image link. This will bring them to an HTML form allowing the missing image to be uploaded via the user's browser. Other types of documents can be uploaded in a similar way.

Special Pages

Wikis are a specialized form of content-management system (CMS). As such, they are programs whose goal is to display the most recent version of a generally static document. This contrasts with many other server-side programs, whose content changes for each user and invocation. Of course, MediaWiki must scan each document as it is published, in order to generate the differ-

ent URLs needed for the existing and new links mentioned within the body of the text. But for the most part, MediaWiki specializes in the easy creation and distribution of static content, rather than programmatic pages.

However, MediaWiki's authors recognized that they and others probably will want to create dynamically generated pages, and thus was born the Special: designation. Any page whose name begins with Special: is treated—well, it's treated differently, populating the document body with the output of a PHP function. In order to create our own special page, we need to write a PHP function and then register that function with MediaWiki.

The first step is to create a new MediaWiki extension, putting it in the extensions directory just under the MediaWiki DocumentRoot. (The basic MediaWiki distribution includes this directory, but it is normally empty after a fresh install.) We will create the file extensions/SpecialHello.php, which looks like this:

```
<?php
$wgExtensionFunctions[] = "wfExtensionSpecialHello";

function wfExtensionSpecialHello()
{
    global $wgMessageCache;
    $wgMessageCache->addMessages(array('hello' => 'Hello page'));

    require_once('includes/SpecialPage.php');
    SpecialPage::addPage(new SpecialPage('Hello'));
}

?>
```

The first line appends our function's name ("wfExtensionSpecialHello") to the global array \$wgExtensionFunctions, putting this function in the directory of extensions. Many MediaWiki extensions change the output from certain tags; in this case, we are looking to create an entire page, rather than modify the behavior of a tag. Nevertheless, we are creating an extension.

Next, we define our function. The first thing that we do is modify \$wgMessageCache, such that our extension will look like a special page rather than a tag-modifying extension. The special page will continue to work if we fail to include these initial two lines, but its listing on Special:Specialpages will look odd, with the title coming between < and > brackets as if it were a tag. Note that the page name must begin with a lowercase letter in the call to addMessages, even when it is capitalized in the call to SpecialPage::addPage. Failing to note this quirk of capitalization will result in strange page titles.

The final two lines import the code specific for special pages, and then create an instance of such a special page, adding it to the directory.

Just what our special page does depends on another file of the same name (that is, SpecialHello.php), located in the includes directory that is parallel to extensions, just under the MediaWiki DocumentRoot. This directory contains a large number of standard special pages that come with MediaWiki, including SpecialNewpages.php, SpecialUserrights.php and SpecialImagelist.php. These functions can access the back-end MySQL database, perform calculations and access external sites—and then pipe the results back into a standard MediaWiki output page.

Here is a simple version of what we might put into our includes/SpecialHello.php file:

```
<?php
function wfSpecialHello() {
    global $wgOut;
    $wgOut->addHTML('Hello, world');
}
?>
```

The above function, which is invoked whenever we go to the Special:Hello

page, adds the HTML "Hello, world" to the output. Notice that we don't have to begin, end or otherwise modify the HTML file that is sent to the user. Nor do we need to worry about choosing a skin, setting up menus or other details. However, we do need to be careful about whitespace, as is often the case with PHP programs—failing to trim whitespace before the initial <?php tag might well produce odd error messages from PHP about modified headers.

Finally, we register our extension and special page in the LocalSettings.php file, adding the following line:

```
require_once("extensions/SpecialHello.php");
```

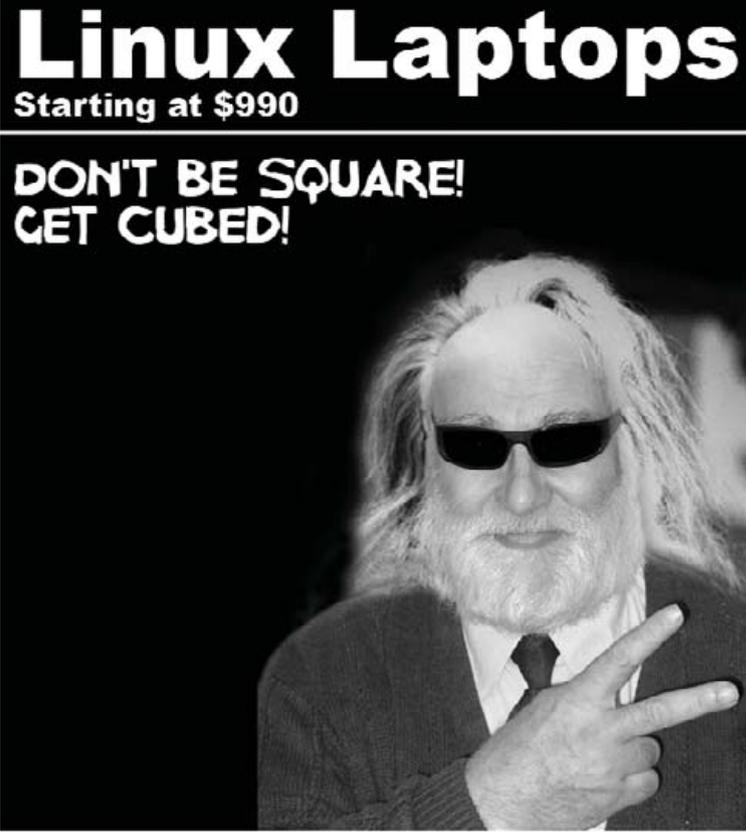
Once you have put the above in place, your site should now have a Special:Hello page, listed as Hello when you visit Special::Specialpages.

Conclusion

Wikipedia has generated a number of headlines over the last few months, and its future as a reliable, neutral, volunteer effort remains to be determined. But Wikipedia has demonstrated the power that a wiki can bring to an organization looking to collect information from a wide variety of participants. If your organization would benefit from a centralized repository created in a decentralized manner, a wiki might just fit the bill.

And, although MediaWiki is far from the only available package, it is highly polished, reliable and easy for both administrators and users. Producing new dynamic pages, known as special pages in the MediaWiki world, requires a fair amount of knowledge of the underlying system. But, getting started is not that difficult, and it is possible (and desirable) to take advantage of the tremendous infrastructure that MediaWiki brings to the table. ■

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.



Linux Laptops
Starting at \$990

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

3 Technologies
Working Hard to Bring Technology to Life

309.34.CUBED
shoprcubed.com

Podcasting for the Penguin!

Jon and Kelly have the Audacity to create high-quality podcasting content on a tight budget.

Jon Watson

Kelly Penguin Girl and I have been producing the weekly GNU/Linux User Show podcast (www.linuxuser.thepodcastnetwork.com) since June 2005. We're just about to hit show #30, which I think makes us one of the oldest running GNU/Linux podcasts aimed specifically at new GNU/Linux users. We've moved on to other things and no longer host the show, but it's been a long, fun journey, and we've learned a lot.

The technical aspects of podcasting on GNU/Linux are on par with podcasting on any other OS, but there are some cultural differences. Podcasting to the Free/Libre/Open-Source Software (FLOSS) community requires a sensitivity of the community's values. In short, a podcaster that intends to podcast to the FLOSS community should really be part of the community. The community wants to be talked with, not talked at. Trespassers will be shot.

We record our shows on either a Kanotix or Linspire box (depending on where we are in the house) using a set of \$25 US NeXXt headsets and two pairs of Y-cables. We thought about getting a mixer, but so far the Y-cables have performed flawlessly, so we don't see the need to spend the extra money on unnecessary gear.

I always find it quizzical why so many podcasters recommend spending \$200 US or more on gear. We've spent less than \$100 US since day one on three sets of headsets and a standalone microphone. One of the lessons that blogging has taught us is that content is king, not the pretty bells and whistles around the content. Although there is certainly a minimum level of quality expected by listeners, the content is what drives the show, not the benefits of shiny microphones and mixers.

Application-wise, we use Audacity to record and edit the show, and EasyTag to insert the ID tags in preparation for file upload.

Audacity

The real workhorse of our show is Audacity. Audacity is a wonderful digital audio workstation (DAW) application that is not only licensed under the GNU GPL and available on SourceForge, but it is also available for all major OS platforms. Audacity has served us well for recording our shows, editing the audio streams, adding effects, importing and aggregating other audio streams and formats, and finally allowing us to export our show in a variety of different formats. Audacity supports Ogg Vorbis encoding out of the box, and it will support MP3 encoding via the LAME encoder (separate download).

One of the killer features of Audacity is the wide range of audio formats it is capable of importing. Over time, we've had to incorporate audio from many sources, such as individual listeners, promo clips from advertisers, audio clips from other shows and downloaded clips from the Internet. Without Audacity's ability to import everything we've thrown at it, we would have been dead in the water many times over.

We generally record the show in many parts. Some are recorded days apart, and some only seconds apart. Regardless, Audacity represents each of these parts as a graphic sound wave, and each of these waves can be manipulated individually (Figure 1).

This intuitive feature makes the post-production of our shows a snap. The tools we use the most are the time-shifter, which allows movement of individual audio parts to snug them up together and kill any dead space, and the insert silence tool. The insert silence tool may sound innocuous, but it's very handy for extending little dead spots to fit around another sound clip.

Audacity also features a very complete set of options that allow granular control of the final audio file quality and size (Figure 2).

Podcasts generally are mostly speech and can therefore be exported at low quality without any appreciable degradation. The Podcast Network's standard is 48Kbps and 22KHz. We used to put out only MP3s of our show, but after Richard M. Stallman came on lucky show #13 and asked us also to produce Ogg Vorbis files, we started doing that as well. It took me a while to understand the Ogg Vorbis compression technique, and our first few Oggs were twice as big as the same show's MP3 file. That didn't make us very popular, let me tell you! Here's the secret: in Audacity, there is no bit rate setting for Ogg Vorbis files. Rather, there is a slider from 0 to 10. A setting of 0, although counter-intuitive, creates a perfectly usable Ogg Vorbis file.

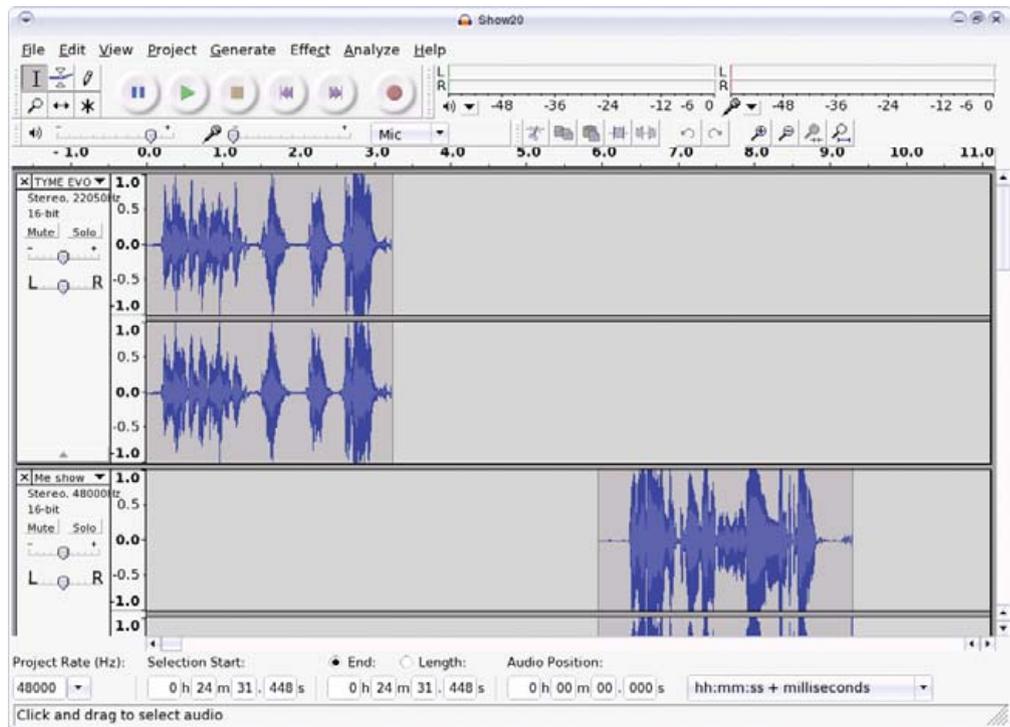


Figure 1. Audacity is the workhorse of this podcast.

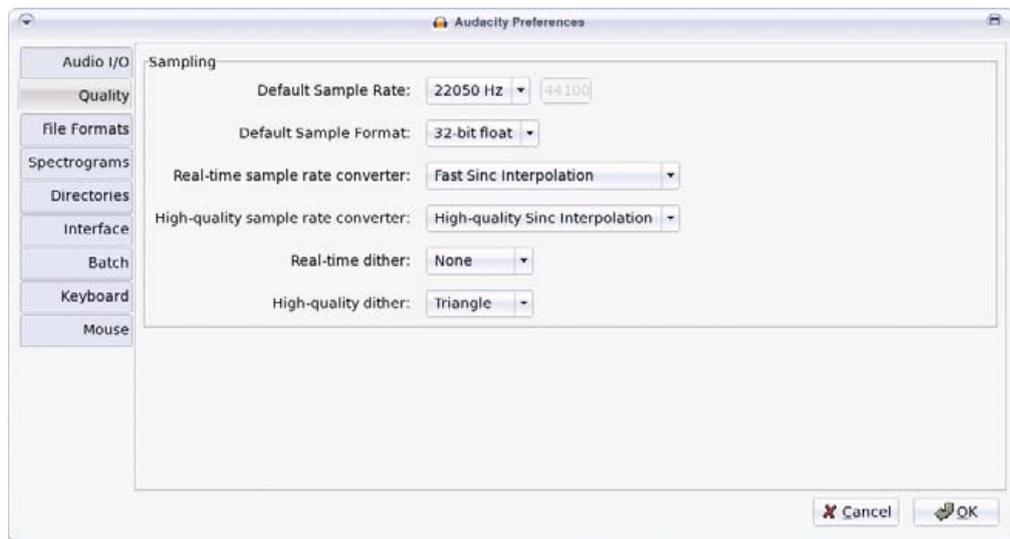


Figure 2. Speech doesn't need the 44KHz sample rate associated with music.

EasyTag

Once the show is created and exported, it's critical to put the ID tags into the file. The ID tags provide the information that scrolls across the screen of listeners' digital audio players (DAP) or their digital audio applications. Without ID tags, listeners would be hard pressed to figure out what show and episode they're listening to. This information isn't necessary only for logistics, it's also critical for promoting your podcast. Listeners can't come back to find your next show if they don't know what the heck they're

listening to.

Audacity has the ability to manage ID tags, but supports only a few fields. The Podcast Network standards required us to supply data for more than those few fields. Therefore, we had to turn to an external tagging application.

We went through a few different tagging applications and finally settled on EasyTag. EasyTag is a nice application that does one thing and does it well. It tags the heck out of Ogg Vorbis and MP3 files. EasyTag is GNU GPL'd and also available from SourceForge.

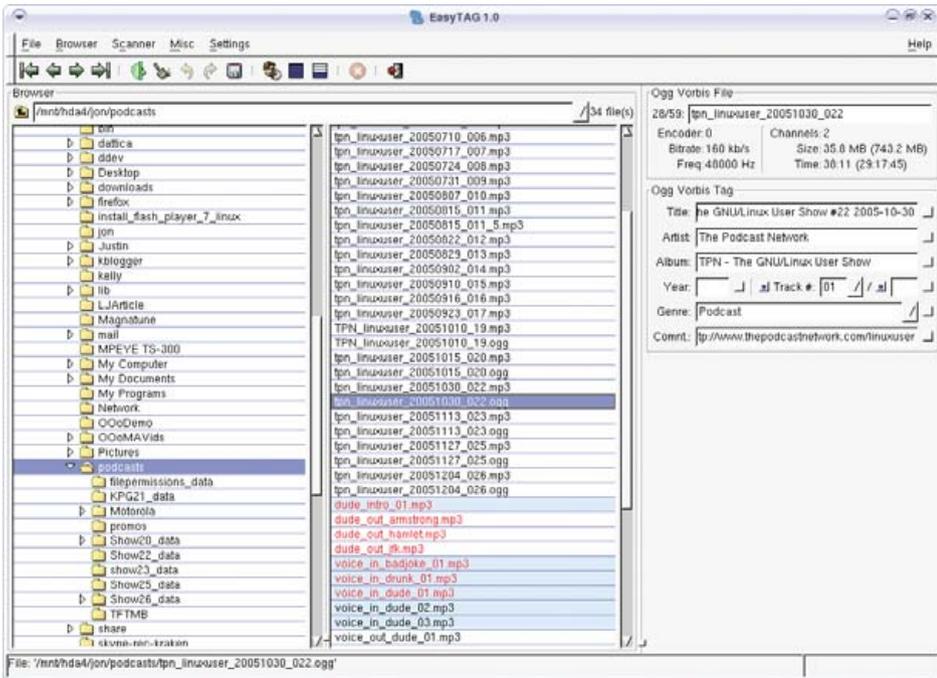


Figure 3. The name EasyTag is appropriate as it makes tagging our MP3 and Ogg files so easy.

There's a real science to tagging podcasts and EasyTag has many more features than we use. Along with the basic functionality of embedding IDv2 and IDv3 tags in both MP3 and Ogg Vorbis files, EasyTag can be set to scan entire directories of audio files and auto-fill in the tags. Because we produce a podcast only once a week, we don't have a lot of use for these advanced features. But, if I had a hard drive full of nontagged music files, EasyTag's scanning feature would be very, very useful.

Aside from the technical aspects of tagging files, there are many different schools of thought on what information should go in to each tag. Most audio players scroll the title, artist and length of the file at a minimum across the screen while playing. Although the title and artist are generally pretty easy to figure out, the title tag requires more thought. Some podcasters put the name of the show and the date it was produced into the title tag. Others feel that the sequence number of show is more important than the date. Both sides typically argue that it's easier for a listener to keep track of a (sequence number or date) than a (date or sequence number).

Podcatchers

This argument was likely more important in the beginning of podcasting, because podcatching software wasn't as advanced as it is now. What makes a podcast a podcast is that it is delivered via an RSS feed. An MP3 (or Ogg file) that is just linked to download on a Web page is just an audio file on the Web, not a podcast. Podcatcher is the affectionate name given to the genre of software that listeners can use to subscribe to these RSS feeds. Once you subscribe to a podcast, the podcatcher should check each show for new episodes and download them automatically. There are varying degrees of complexity in today's podcatchers, but most offer at least the check and download new episodes functionality.

One of the more popular GNU/Linux podcatchers is BashPodder [see Marcel Gagné's article on page 32 for more information on BashPodder] written by Linc Fessenden of The Linux Link Tech Show. Along with the basic BashPodder, Linc also wrote BPGUI, which is a nice GUI front end for the command-line BashPodder client. In true community fashion, Linc released BashPodder under the GNU GPL, and many people have made modifications to the base application. A quick Google search for the term BashPodder shows the wide variety of improvements and changes the community has made to it. Whatever your taste, it's likely that you will be able to find a flavour of BashPodder that meets your needs.

The stable of podcatchers for GNU/Linux is growing as podcasting becomes more popular. CastPodder [see Marcel Gagné's article on page 32 for more on CastPodder] is another popular podcatcher, and even amarok has podcatching capabilities.

I cannot stress enough that content is what listeners tune in for. Audio quality is important, but it's not the Holy Grail. Good guests, solid content, credible hosts and regular production are what build an audience.

Happy podcasting! ■

Jon Watson is the host of the weekly GNU/Linux User Show on The Podcast Network. Jon has written articles for *Really Linux*, *Linux Journal*, has been interviewed on the topic of podcasting for *Alberta Venture Magazine* and is slated to speak at the Calgary Linux User Group Linuxfest in spring of 2006. In his spare time, Jon also writes the *New Linux User* (www.newlinuxuser.com) blog for b5 Media (www.b5media.com) and can be contacted at me@jonwatson.ca. Jon lives with his fiancée and co-host Kelly Penguin Girl in mountainous Calgary, Alberta, Canada.

4th Annual 2006 Linux/Open Source on Wall Street show and Conference

APRIL 24, 2006 (Mon) ROOSEVELT HOTEL, NYC
Madison Ave and 45th, next to Grand Central Station



The Largest New York Meeting for Linux and Open Source for Wall Street.

The 4th Annual Linux/Open Source on Wall Street will assemble Wall Street and global financial services IT decision-makers for a focused show and conference for Linux and Open Source and other tools and techniques for the Linux marketplace. Wall Street IT managers, project managers and Linux and Open Source decision-makers will assemble at this New York marketplace to see the latest solutions to lower costs and total cost of ownership. The show will provide face-to-face contact with exhibitors and 1,000 attendees at the convenient Roosevelt Hotel at Grand Central Station.

Show Hours: Mon, April 24 8 - 4:00 pm
Conference Hours: 9 - 4:50 pm

Register today for your free Show badge at:
www.linuxonwallstreet.com

LINUX ON WALL STREET

For more information on the event, to be a sponsor or exhibitor, contact:

Flagg Management Inc
353 Lexington Ave, New York, NY 10016
212-286-0333 Fax: 212-286-0086
flaggmgt@msn.com

2005 Gold Sponsors



O'REILLY®

JULY 24-28, 2006 • PORTLAND, OREGON

EIGHTH ANNUAL
OSCON
OPEN SOURCE
CONVENTION

OPENING INNOVATION

OSCON 2006 will feature the projects, technologies, and skills that you need to write and deploy killer modern open source apps.



Linux

Perl Conference 10

PHP Conference 6

Python 14 Conference

Apache

XML

Emerging Topics

Java

Ruby on Rails

Security

Databases

REGISTER BY JUNE 5 AND SAVE!

conferences.oreilly.com/oscon

Creating a Planet Me Blog Aggregator



Build a personal blog aggregator using the same code that powers many “Planets” sites, such as Planet Apache.

Ben Martin

The Planet Project allows on-line communities to build a central Web page easily, aggregating blogs from people in their community. The Planet code powers such community blogs as Planet GNOME and Planet Apache. Such on-line uses of the Planet code provide a low cost of entry for people to keep an eye on a community. This article focuses on using the Planet code on your local machine to create your own custom blog aggregator.

Installation

The Planet code requires Python 2.2 or later. The simplest method to install Planet is to download a nightly snapshot tarball from the planetplanet.org Web site and extract it to your home directory. I tend to rename the extracted planet-nightly directory to include its day of download and use a handy link to the current version of Planet Me.

In this article, I've used references to the path of my home directory a few times; remember to substitute your own home directory in the examples.

The two final commands in Listing 1 show how to fetch current news feeds and set up your initial Planet. The commands will vary depending on things such as whether or not you have to use a proxy server to access the Internet. After running these commands, you should have a Planet Me viewable in your Web browser at `~/planet/me/index.html`. After doing these steps, your planet should look similar to Figure 1.

You'll want to customize which news feeds you are viewing. This is done at the end of `me-meta/config.ini`. The configuration file defines a section by text surrounded by square brackets. Options for a section follow its initial definition as key=value pairs. You define each blog to aggregate in a section where you specify the URL of the RSS feed for the section name. See Listing 2 for an example from the default `config.ini` file.

The name will be shown in the header for each aggregated post from that blog, and the face image will be on the right side when using the default HTML templates. The `facewidth` and `faceheight` are optional by default.



Listing 1. Installing Planet

```
$ cd ~
$ tar xjvf planet-nightly.tar.bz2
$ planetdated=planet-$(date +%d%b%y')
$ mv planet-nightly $planetdated;
$ ln -s $planetdated planet
$ cd planet
$ cp -av fancy-examples me-meta
$ cd me-meta
$ cp ../examples/*.xml* .
$ edit config.ini
name = Planet Me
link = file://home/ben/planet/me/index.html
owner_name = John Doe
owner_email = root@localhost

# later in the file
# template_files should all be on one line
template_files = me-meta/index.html.tpl
me-meta/rss20.xml.tpl me-meta/rss10.xml.tpl
me-meta/opml.xml.tpl me-meta/foafroll.xml.tpl

# later in the file change
# fancy-examples/index.html.tpl
[me-meta/index.html.tpl]
items_per_page = 30

$ cd ..
$ mkdir cache
$ ln -s output me

# Without proxy
$ python planet.py me-meta/config.ini

# Using a standard squid proxy on "dairiserver"
$ http_proxy=http://dairiserver:3128/ \
python planet.py me-meta/config.ini
```

Listing 2. Sample Aggregation Definition

```
[http://www.gnome.org/~jdub/blog/?flav=rss]
name = Jeff Waugh
face = jdub.png
facewidth = 70
faceheight = 74
```

Listing 3. How to Get the Image from Slashdot

```
$ cd ~/planet/me/images/
$ wget \
http://images.slashdot.org/topics/topicslashback.gif

# convert is from ImageMagick
$ convert topicslashback.gif slashdot.png
```

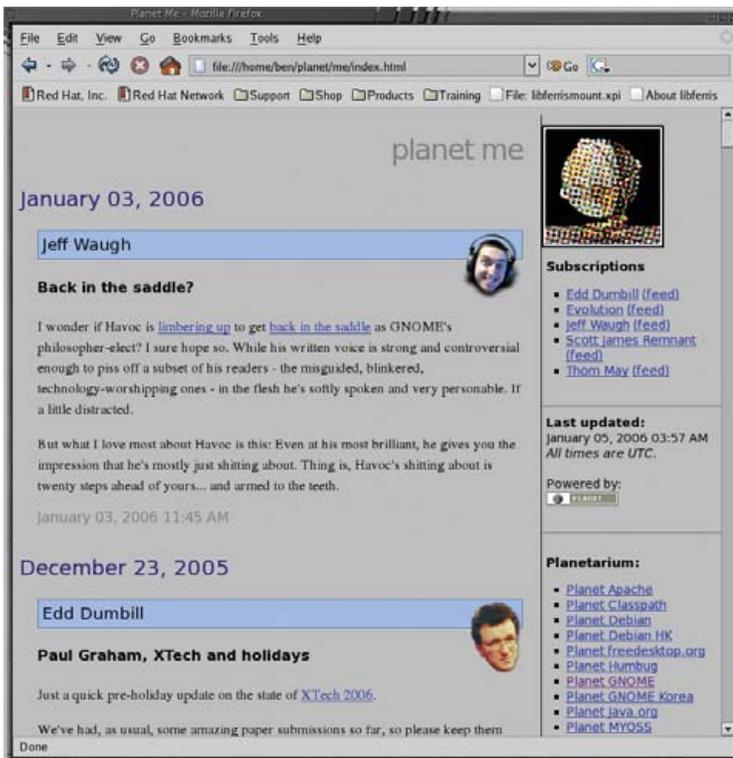


Figure 1. A New Running Planet Installation

Listing 4.

Create a Definition to Use the Slashdot Icon

```
$ edit ~/planet/me-meta/config.ini
[http://rss.slashdot.org/Slashdot/slashdot]
name = Slashdot
face = slashdot.png

$ cd ~/planet
$ python planet.py me-meta/config.ini
```

Listing 5.

Set Up a cron Job to Aggregate Blogs

```
$ mkdir -p ~/mycron
$ cd ~/mycron
$ vi upd-planet.sh
#!/bin/sh
cd ~/planet;
http_proxy=http://dairiserver:3128/ \
python planet.py me-meta/config.ini
$ chmod +x upd-planet.sh
$ echo \
'00 04 * * * /home/ben/mycron/upd-planet.sh' \
>|upd-planet.cron

# only if you already use cron from outside ~/mycron
$ crontab -l >|oldcrontab.cron

$ cat *.cron >|newtab
$ crontab newtab
$ rm -f oldcrontab.cron
```

Listing 6.

Generate files to define the blogs to aggregate.

```
$ cd ~/planet/me-meta
$ mv config.ini config.ini.template
$ edit config.ini.template
# remove all blog URL sections from the bottom of file
# search for http: to find the first one

$ mkdir blogs
$ echo http://rss.slashdot.org/Slashdot/slashdot \
>blogs/slashdot.blog
$ ./generate-config
```

Many sites provide handy topic icons that can be used to spruce up your Planet Me. For example, in Listing 3, I use one of the Slashdot section icons (see the on-line Resources) for news items taken from Slashdot's RSS feed.

Assuming you use the Planet setup as described in this article, the topic icons are stored in ~/planet/me/images. You can see the setup for my Slashdot topic icon in Listing 3.

Listing 4 shows the new section to append to the config.ini to integrate the Slashdot icon into your Planet Me.

Listing 7.

Use the files to create a blog aggregation configuration.

```
#!/bin/sh

cp -av config.ini.template config.ini
for if in blogs/*.blog
do
    base=$(basename $if .blog);
    content=$(cat $if);
    echo "" >> config.ini
    echo "[$content]" >> config.ini
    echo "name = $base" >> config.ini
    echo "face = $base.png" >> config.ini
done
```

Listing 8.

Set conditions for your output with the TMPL_IF tag.

```
<TMPL_IF title>
<a href="<TMPL_VAR link ESCAPE="HTML">">
<TMPL_VAR title>
</a>
</TMPL_IF>
```

Dynamic Content

Now you need to have the Planet code run to aggregate blogs, and make it easy to modify the list of blogs to aggregate.

You can solve the regular aggregation by using cron. Listing 5 shows how to have Planet Me updated every night.

It is easy to add and remove blogs if you keep a list of blog definition files instead of trying to manage the configuration file itself manually. You can use the generate-config script shown in Listing 6 to move the blog name and URLs into very simple files in a blog subdirectory.

You can use a file manager or the command line to add or remove files that determine the configuration of your aggregation scheme. This also paves the way for a simple Firefox extension to allow new RSS feeds to be added to Planet Me from a context menu. Handling archives as shown later is also simplified by moving the blog information out of config.ini.

Updating the Look and Feel

The two files that control how your planet will look are me-meta/index.html.tpl, which is the template for the page content, and me/planet.css, which is the cascading stylesheet.

By default, the face, entry, date and sidebar all define styles that can be changed using the stylesheet. You can use custom fonts by modifying the font-family CSS tag.

The index.html.tpl template has extra tags that the Planet code uses to generate the final index.html file. The main tags of interest are TMPL_LOOP, TMPL_IF and TMPL_VAR. The news feeds are placed into the output page using the <TMPL_LOOP Items> HTML-like tag and its corresponding close tag. The HTML elements between these two tags will be output once for each news item to be displayed. These elements define what and how output is generated for each news item.

The Planet code uses these variables to get at the news feed

Listing 9.

New Channel Items Section for ~/planet/me-meta/index.html.tpl

```
<TMPL_LOOP Items>
<TMPL_IF new_date>
<h2><TMPL_VAR new_date></h2>
</TMPL_IF>

<div class="news-item-icon">
  <a href="<TMPL_VAR channel_link ESCAPE="HTML">"
    title="<TMPL_VAR channel_title ESCAPE="HTML">">
    <TMPL_VAR channel_name>
  <br/>
  " >
  </a>
</div>

<div class="news-item">
<TMPL_IF title>
<h4><a href="<TMPL_VAR link ESCAPE="HTML">">
  <TMPL_VAR title></a></h4>
</TMPL_IF>
<div class="entry">
<p>
<TMPL_VAR content>
</p>
<p class="date">
  <a href="<TMPL_VAR link ESCAPE="HTML">">
<TMPL_IF creator>by <TMPL_VAR creator> at </TMPL_IF>
<TMPL_VAR date></a>
</p>
</div>
</div>

</TMPL_LOOP>
```

content. For example, it replaces the <TMPL_VAR title> tag with the actual title of the current news item. Note that TMPL_VAR doesn't have a corresponding close tag.

The TMPL_IF tag is used to check whether information exists or to set specific conditions. For example, sometimes news items do not have title information. The code in Listing 8 will output title information if it exists, and output nothing if a title does not exist. The escape attribute on the TMPL_VAR tag tells Planet to make sure that the value of the link variable is in a form that is a legal HTML attribute.

You'll have to edit both the me-meta/index.html.tpl and CSS files to move the channel icon to the left of the news item with Planet Me.

By default, your index.html.tpl will display the channel icon only when the current news item is from a different channel than the one preceding it.

I've removed the <TMPL_IF new_channel> tags from around the outputting of the face image information in the fragment of index.html.tpl shown in Listing 9. I also used a CSS class of news-item-icon for the channel image and news-item for the main news post section and a new class of embedded-face for the actual channel image.

Shown in Listing 10 is the new stylesheet code to set the channel image on the left of the news item. Your Planet Me should now appear



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



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 10.

New Styles to Add to ~/planet/me/planet.css

```
div.news-item-icon {
    float: left;
    position: relative;
    left: 4px;
    margin-top: 25px;
    padding: 0 20px 30px 0;
    width: 120px;
    text-align: center;
}

div.news-item-icon a {
    text-decoration: none;
}

div.news-item {
    margin-left: 140px;
}
```

Listing 12.

How to Archive News Feeds on a Repeated Basis

```
$ cd ~
$ unzip Jena-2.3.zip
$ edit ~/.bashrc
# append a handy classpath setup
JenaSetup() {
    for if in ~/Jena-2.3/lib/*.jar; do
        export CLASSPATH=$CLASSPATH:$if;
    done
}
$ . ~/.bashrc
$ JenaSetup

# archive news feed
# repeatable three step
$ cd ~/planet/me
$ mv -f archive.xml rss10-archive.xml
$ java jena.rdfcat rss10*.xml >archive.xml
```

Listing 11.

Another Way to Use TEMPL_IF

```
<TEMPL_IF channel_foo>
Have foo:<TEMPL_VAR channel_foo ESCAPE="HTML">
</TEMPL_IF>
```

like Figure 2. If you prefer your channel icons to the right of the news item, change the stylesheet to have news-item-icon's float tag be right and news-item's margin-left be 0px.

The face=whatever.png line used in the channel definition files is not special to Planet. You can define any other variables you want on a per-channel basis, and they will be available for use in your index.html.tmpl. For example, Listing 11 shows the use of an optional variable foo, which might be defined for a channel as foo=bar after a channel description in your config.ini file.

Another great way to learn how to customize your Planet Me is by visiting other Planet Web sites. Examine their HTML and CSS files to learn how they modified the look and feel.

Keeping and Viewing Archives

The Planet code was designed to aggregate news feeds from many sources and supply a recent history of them on a single page. For local use with Planet Me, it is nice to be able to see a news feed for an arbitrary period in the past.

Your Planet Me will create a valid RSS RDF news feed that you can use to archive your Planet. All things in an RDF file revolve around triples. The three parts of a triple are referred to as the subject, predicate and object. An example triple might say that a news item has a given publication date, for example, item57 has-date 3-Jan-2006. An RSS news feed defines a news channel, associates that channel with a list of news items and defines interesting properties for each news item, such as its title, publication date and text content. Usually things like has-date are defined using long URIs to avoid two triples accidentally having the same literal value.

An easy, yet powerful way to archive your Planet's RSS is to use the Jena Project. Once you have a Java virtual machine installed, all you need to install Jena is to download a tarball, extract it and add it to your classpath.

Shown in Listing 12 are the installation steps and repeatable archiving process for news feeds. You could place your news feed archive into a database using Jena if you are collecting many feeds over a long time. Jena gives you the ability to use very powerful queries against your archive to re-create your Planet.



Figure 2. My Sample Planet Site with a Customized Template and CSS File

Listing 13.

Use Jena to query your Planet Me site.

```
$ cat rss-by-date.sparql
PREFIX dc: <http://purl.org/dc/elements/1.1/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

DESCRIBE ?channel ?bnode ?a WHERE
{
  ?channel ?items ?bnode .
  ?bnode ?hasitem ?a .
  ?a dc:date ?date .
  FILTER ( xsd:dateTime(?date)
    >= xsd:dateTime("2006-01-03T00:00:00")
    && xsd:dateTime(?date)
    <= xsd:dateTime("2006-01-05T00:00:00") )
}

$ cd ~/planet/me
$ java jena.sparql --data archive.xml \
--query rss-by-date.sparql --results RDF/XML \
>my-query-result.rss
```

Listing 14.

Use Jena query results to modify your Planet Me aggregation.

```
$ cd ~/planet/me-meta
$ cp -av config.ini.template config.ini
$ echo \
"[file:///home/ben/planet/me/my-query-result.rss]" \
  >>config.ini
$ echo "name = archive" >>config.ini
$ cd ~/planet
$ rm -f cache/file.home*
$ python planet.py me-meta/config.ini
```

Listing 15.

You can use regular expressions as filters.

```
PREFIX dc: <http://purl.org/dc/elements/1.1/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX rss: <http://purl.org/rss/1.0/>
PREFIX content: <http://purl.org/rss/1.0/modules/content/>

DESCRIBE ?channel ?bnode ?a WHERE
{
  ?channel ?items ?bnode .
  ?bnode ?hasitem ?a .
  ?a content:encoded ?content .
  FILTER ( regex(?content, ".*product.*", "i") )
}
```

Listing 13 shows a simple time interval query for news feeds. This query is in the SPARQL query language, which is used for querying RDF repositories. The core of the query links the channel, news item and date components before applying a filter to which news items are to be returned based on the date attached to that news item.

We can now easily change Planet Me to use only your query results as input, as shown in Listing 14, because we moved the blog URLs and meta-data into separate files as explained above.

The channel icons will all be the same using the above query, because you are querying a single news feed, your own. Another expression that can be used in the FILTER section is regex(). The example shown in Listing 15 filters all news items and shows only those that match the case-insensitive regular expression.

Wrap-Up

Planet Me is designed to be used for on-line community blog aggregation, but the Planet code can make a very effective blog aggregator for personal use. The Planet code is designed to create on-line blog aggregates that are viewed by a large amount of people. With some tinkering, the Planet code can make a very effective personal blog aggregator giving you the freedom to choose explicitly who is in your community as well as easily creating archives of your Planet and searching past news using a very powerful query language.

Resources for this article: www.linuxjournal.com/article/8830. ■

Ben Martin spends most of his time working on virtual filesystems and data mining over them. Recent joys include extending libferris to allow mounting Emacs and Firefox as filesystems.

Micro-Server Internet Appliance Engine

SOM-5282EM

- Coldfire 66/80 Mhz CPU
- 10/100BaseT Fast Ethernet
- 16 GP I/Os & 8 Channel A/D
- Programmable in Java™ or C
- Telnet, FTP, and HTTP Servers
- 3 Serial Ports, CAN 2.0B & SPI
- uClinux with Real Time support
- 4.5 MB Flash & up to 16 MB RAM
- Typical Power Consumption < 2 Watts
- Real Time Clock & Nonvolatile Memory
- Carrier/Socket Board & Power Supply Available
- Small, 144 pin SoDIMM form factor (2.66" x 1.5")

The SoM-5282EM is a System on a Module, based on the Freescale MCF5282 Processor. This 32-Bit processor runs uClinux making it extremely easy to create smart Network/Internet capable devices, with Data Acquisition and Control properties. If Real-Time processing is required we can optionally provide RTAI Real-Time extensions. Write sophisticated network applications in days instead of months using standard GNU tools. Single Unit pricing starts at \$150. Optional Carrier/Socket board, Enclosure, and Power Supply also available. For additional information visit, www.emacinc.com/som/.

Since 1985
OVER
20
YEARS OF
SINGLE BOARD
SOLUTIONS

EMAC, inc.

EQUIPMENT MONITOR AND CONTROL

Phone: (618) 529-4525 • Fax: (618) 457-0110 • Web: www.emacinc.com



MOVING TO BLOG-CITY

**BLOG-CITY.COM IS A
LINUX-BASED FEATURE-RICH
BLOGGING COMMUNITY THAT
DELIVERS A ROBUST AND
HEAVILY CUSTOMIZABLE
BLOGGING EXPERIENCE.**

DEE-ANN LEBLANC



Figure 1. The Blog-City Home Page

One combination free/pay blog site available on the Internet is Blog-City.com. This site runs on a farm of Dell PowerEdge Blade servers, each with six dual-CPU blades running Red Hat Linux 9. Blog-City.com runs Java's JDK 1.4.x underneath BlueDragon (www.newatlanta.com/products/bluedragon). BlueDragon is the same ColdFusion Markup Language (CFML) application that runs MySpace.com. MySpace.com is the fifth largest Web site in the world.

Blog-City offers no static content. It dynamically constructs every page as it is requested by pulling it from a MySQL 4.1.x database that is replicated for both backup and administrative purposes. Blog-City uses MySQL 5.0 for a caching server to speed up access time, and it uses MySQL 4.1.x for all of the access and referrer statistics. These separate database servers keep the load for particular types of access from bogging down other types of access.

For Web service, Apache 1.3.x communicates with BlueDragon through mod_backhand (www.backhand.org/mod_backhand/) with a Java library the company wrote and then released back to the community to make sure that the load balancing works properly for the application server. The mod_backhand module load balances only the Web server by default. Blog-City uses additional mini-farms to produce RSS feeds and answer search engine queries. There is also custom software for handling spam, denial-of-service attacks and a bug-tracking system. All in all, Blog-City is an operation that is spread over a number of time zones and multiple continents, just like its users.

Creating a Blog

Before trying to create a blog in Blog-City, try various URLs in the format of word.blog-city.com to see which ones are already in use (for example, kleekai.blog-city.com). Once you find something you think is available, go to www.blog-city.com (Figure 1) and find the Free Signup box.

Enter the name you want to use for your blog in the text box. For the above example, you would enter kleekai. Next, click Submit to open the Create New Blog dialog box (Figure 2) and walk through the rest of the blog creation process.

Go to your blog's URL to log in to your blog for the first time. You also



Figure 2. The Blog-City Create New Blog Dialog Box

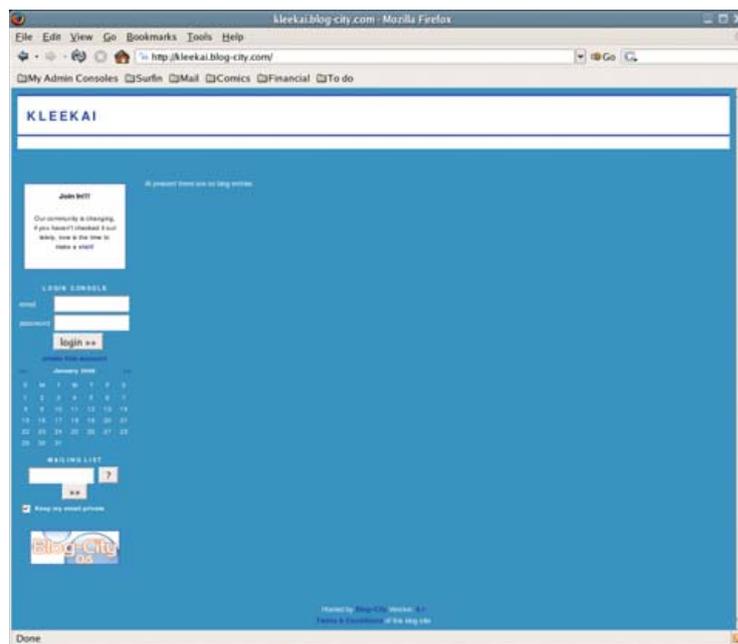


Figure 3. The Initial Klee Kai Blog Page

can click through one of the links on the creation page or click one of the links in your confirmation e-mail. When you do, you will see your initial blog page (Figure 3).

Enter your e-mail address and password in the Login Console text boxes, and then click Login. Once you reach the Main Administration page (Figure 4), bookmark it. This is where you need to go whenever you want to change something on your blog or check your statistics. Much like UNIX, there is often more than one way to get to a particular option through this interface. This article focuses on one path for each.

The tabs across the top of this page lead to:

- Home: view Main Administration page.
- Blog-It: create a new blog entry.

- **Entries:** view and delete existing posts, and see how many times each item has been read and commented on. Blog-City supports a variety of entry types, including regular blog posts, reviews, photo albums, events and podcasts.
- **Look'n'Feel:** control your blog's layout, themes, headers, menus and more.
- **Settings:** alter your personal and blog settings.
- **Extra:** access statistics, account details and more.

Along the left, what you see depends on whether you have a premium (paid) account or not. The top grouping of links leads to:

- **Home:** the Main Administration page.
- **Write a new blog entry:** the blog entry creation page.
- **View Latest Comments:** the most recent comments made to your blog, the entry to which they were made, when they were made and the ability to delete, reply or report them as spam.
- **View Latest Trackbacks:** premium-only—rather than replying in comments, some people create trackbacks, which are essentially replies made on their own blogs. By clicking this link and then Change Your Trackback Options, you can tell Blog-City to send you an e-mail whenever a trackback is made.
- **View Latest Statistics:** premium-only—takes you to your Statistics page, where you can view the most-read blog posts, the most-commented, how many hits you have had per month and within a month, the browsers that have been used to access your blog, what pages referred people to your blog and which are your top podcasts.
- **View Podcast Statistics:** premium-only—see which are your most popular podcasts.
- **Manage Your LinkBlog:** a LinkBlog is a list of links with comments displayed for each. This link lets you create a LinkBlog and manage its contents. It also provides the URL and RSS feeds where you and others can find the content.

Beneath this group on the left are:

- **Hits since date:** date is either when the blog was created or the last time you reset the hit counter. To reset the counter, click the date and then click OK. This action does not reset any other statistics.
- **Total Blog Entries:** how many posts you've made on your blog, and a statistic of how many posts that makes per day.
- **Total Comments:** how many comments people have posted.
- **Mailing-List Subscribers:** how many people have signed up to receive e-mail notifications when you put up a new post. Click through the link to see non-anonymous subscribers.

There are many more items in the Main Administration page, including:

- **Recently Published Entries:** the last five blog entries you published.

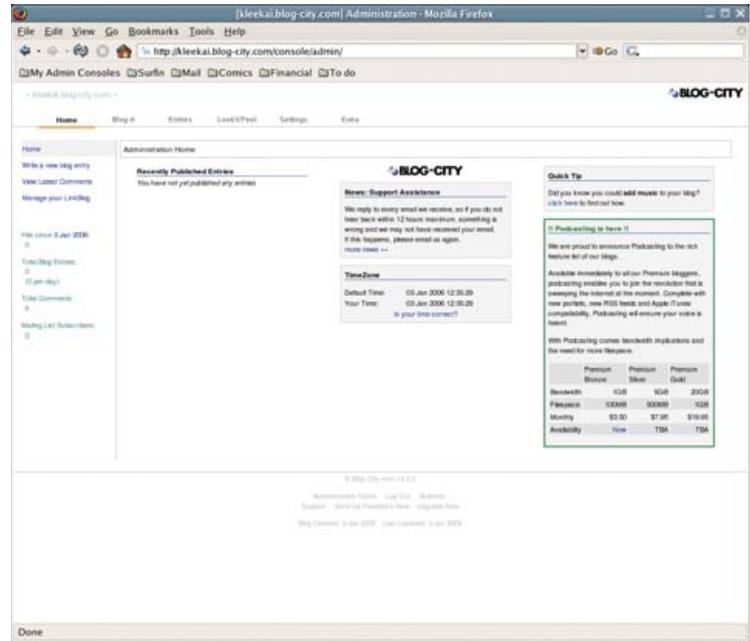


Figure 4. The Blog-City Main Administration Page

- **Entries not yet published:** entries you saved as drafts rather than publishing.
- **My Notes:** any notes you have created (go to Extra→Your Notes to do so).
- **News:** the most recent News posting from the Blog-City staff.
- **TimeZone:** your local time for you to check.
- **Quick Tip:** a tip that may tell you how to do something you didn't know you could do.
- **Additional information:** includes pricing and availability of new features, such as podcasting.

You'll find additional options related to RSS, posting by e-mail and more by clicking the Settings tab. Also, note the link to your own blog in the upper-left corner. At any time, you can click this link to view your own blog. Clicking through here makes sure that your test views aren't counted among your viewer statistics, which is a nice way to make sure they don't get artificially inflated.

It's All about Looks

Most people want to get their blog looking the way they want it before they start to enter content. To do this, click the Look'n'Feel tab to change to the Main Blog Layout page (Figure 5). Here, you see a visual representation of your layout with labels in place. Those that are clickable are items you can change. Those that aren't are for premium accounts.

The links along the left of this page are:

- **Main Blog Layout:** the layout chart you're looking at now.
- **Custom Home Page:** premium-only—lets you design a new layout chart completely from scratch.

- Layout Style—lets you choose how the page should be laid out from a list of templates. Premium users get many more layout options.
- Theme Editor: lets you change your page's theme.
- Menu Editor: premium-only—lets you create menus to appear on your site that can be used for navigation purposes.
- Bookmarks: lets you create bookmarked links and organize them.
- Visitor Polls: premium-only—lets you create polls to be added as portlets.
- News Feeds: premium-only—lets you set what RSS feeds to subscribe to, and then display it in a subset page (such as kleekai.blog-city.com/newsfeed), or in a portlet.
- Custom Portlets: premium-only—lets you create and manage portlets of your own using the HTML editor window.
- Blog Board: lets you access your Blog Board (portlet where users can leave quick messages or chat) content, clear it out or set delays on when comments will appear.
- Neighborhood: premium-only—lets you create a subpage (such as kleekai.blog-city.com/neighborhood) that contains a list of other blogs you track. You can mark this page to be linked off of your main page.

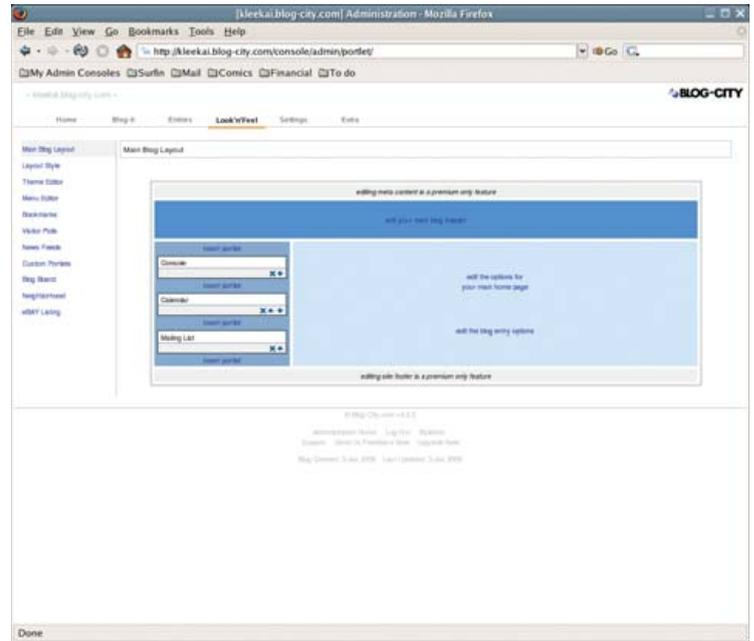


Figure 5. The Blog-City Main Blog Layout Page

- eBay Listing: premium-only—lets you show the latest items for sale in

Don't complicate a simple task

Keep basic tasks just that with handheld, stationary and vehicle-mounted wireless data collection from AML. While others are busy reinventing the wheel, we're keeping things simple, from our products to our personal service. Visit us at www.amltd.com or call **1.800.648.4452** for a real live person.

M7100 Wireless Family

Advertiser Index

For advertising information, please contact our sales department at 206-782-7733 ext. 2 or ads@ssc.com.
www.linuxjournal.com/advertising

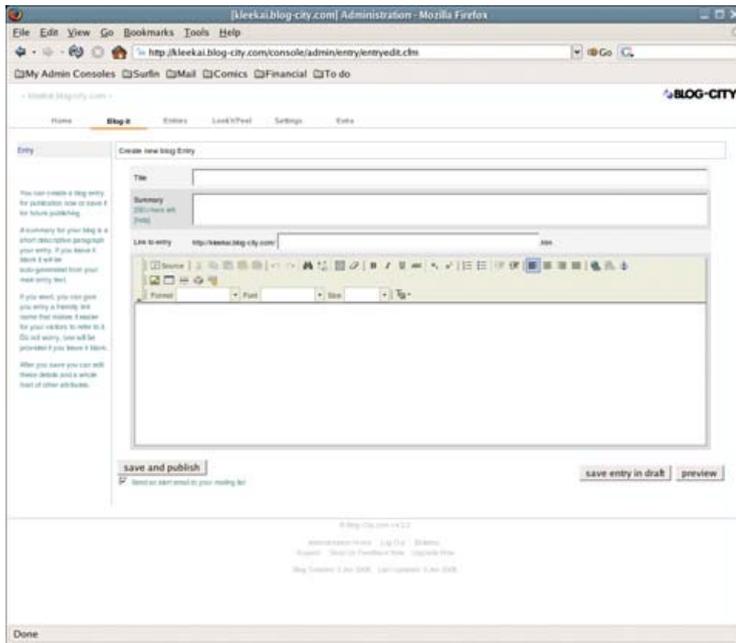


Figure 8. The Blog-City Create New Blog Entry Page

New Blog Entry page (Figure 8), and then fill in the fields:

- Title: the title for the post.
- Summary: you either can fill this in explicitly or let the software do it for you by pulling out the first 250 characters from your post. As you might imagine, this can end up trailing off in the middle of a sentence.
- Link to entry: lets you tell Blog-City to use a specific URL instead of using the title as the URL, which is handy if you need to be able to spell out the link to people later.
- The content box: enter the content for your blog.

When you're finished, click Save and publish if you're finished with the post and ready to put it up, click Save entry in draft if you're not ready to publish it, or click Preview if you want to see what the post will look like. If you have people signed up for your mailing list (look in Extra→Mailing List to see how many members—people can sign up right from your front page), you want to have checked the Send an alert e-mail to your mailing-list check box before posting. If you saved the entry to draft, you can post it later by clicking the Home tab, looking in the Entries not yet published section and clicking the entry title, and then clicking Publish Entry to the left of the content.

Click the Entries tab to add other types of content. An Entry generally refers to a blog post, but you also can post reviews, photos and Events. The administration section changes as often as the Blog-City staff adds new features, so be sure to explore it as time goes on, so you don't miss something new and cool.

Happy blogging! ■

Dee-Ann LeBlanc (dee-ann.blog-city.com) is an award-winning technical writer and journalist specializing in Linux and miniature huskies. She welcomes comments sent to dee@renaissoft.com.

Advertiser	Page #	Advertiser	Page #
ABERDEEN, LLC www.aberdeenc.com	37	LINUX JOURNAL www.linuxjournal.com	94
ADVANCED CLUSTERING TECHNOLOGIES www.advancedclustering.com	25	LINUX ON WALL ST. www.linuxonwallstreet.com	68
AML www.amltd.com	79	LPI www.lpi.org	55
APPRO HPC SOLUTIONS appro.com	C2	MBX www.mbx.com	9
ASA COMPUTERS www.asacomputers.com	73, 83	MICROWAY, INC. www.microway.com	C4, 53
BITDEFENDER, LLC www.bitdefender.com	51	MIKRO TIK www.routerboard.com	7
CARINET www.cari.net	91	OPEN SOURCE STORAGE www.opensourcestorage.com	93
CIARA TECHNOLOGY www.ciara-tech.com	10, 11, 16, 17	O'REILLY MYSQL USERS CONFERENCE www.mysqluc.com	69
CONCURRENT COMPUTER CORPORATION www.ccur.com	30	PENGUIN COMPUTING www.penguincomputing.com	61, 63
CORAID, INC. www.coraid.com	35	POLYWELL COMPUTERS, INC. www.polywell.com	41
COYOTE POINT www.coyotepoint.com	3	THE PORTLAND GROUP www.pggroup.com	28, 29
CYCLADES CORPORATION www.cyclades.com	1	RACKSPACE MANAGED HOSTING www.rackspace.com	C3
EMAC, INC. www.emacinc.com	75	R CUBED TECHNOLOGIES www.rcubedtech.com	65
EMPERORLINUX www.emperorlinux.com	15	SBE, INC. www.sbei.com	23
FAIRCOM CORPORATION www.faircom.com	95	SERVERS DIRECT www.serversdirect.com	31
FOURTH GENERATION SOFTWARE SOLUTIONS www.fourthgeneration.com	40	SUGARCRM, INC. www.sugarcrm.com	43
GENSTOR SYSTEMS, INC. www.genstor.com	85	SUPERMICRO www.supermicro.com	13
HURRICANE ELECTRIC www.he.net	45	SWELL SOFTWARE, INC. www.swellsoftware.com	92
IPSCON www.ispcon.com	89	TECHNOLOGIC SYSTEMS www.embeddedx86.com	20
IRON SYSTEMS www.ironsystems.com	57	THINKMATE www.thinkmate.com	19
JAVA ONE java.sun.com/javaone/sf/index.jsp	49	TYAN COMPUTER USA www.tyan.com	5
LANTRONIX www.lantronix.com	21	USENIX www.usenix.org/usenix06/lj	47
LAYER 42 NETWORKS www.layer42.net	87	ZERVEX www.zervex.com	6
LEVANTA www.levanta.com	39	ZT GROUP INTERNATIONAL www.ztgroup.com	33

XOOPS,

YOU CAN DO IT AGAIN AND AGAIN

How to set up, install and configure XOOPS, the object-oriented extensible blog and content management system. **JUAN MARCELO RODRIGUEZ**

XOOPS is a content management system that allows you to handle a wide range of actions and operations in a dynamic and flexible way. XOOPS covers the needs of a large spectrum of users. Whether you are a common user looking for blog tools, or an advanced one who needs more complex or professional solutions, XOOPS fits the bill for both situations.

XOOPS is based on object-oriented PHP code. The source code is designed to make XOOPS extensible. You can write your own modules to extend its feature set, or download optional XOOPS modules that people contribute and share on the Internet.

This article covers the most relevant aspects of this helpful tool and examines common problems and gotchas.

One can quickly notice that XOOPS provides a lighter and simpler setup than similar content management systems, such as Plone and Drupal. Some may argue that Plone and Drupal have more depth, but the simple and clean structure of XOOPS makes it ideal for those who are getting started with content management tools.

I have been working on XOOPS for three years now, with a local Linux User Group. I have found XOOPS, during all this time, to be far from disappointing. Each particular need I've ever had was fulfilled by XOOPS. What's more, being GPL software, XOOPS gave me the chance to write new modules as I pleased or simply to modify the old ones. Needs require solutions, and solutions present challenges, but when challenges are taken by one's own hand, reward pays double.

Speed is another XOOPS advantage. Even when you have a large user database and hundreds of posts, XOOPS response time is still among the best.

Requirements and Setup

If you have ever used a gallery system, a blog or a content management system, you are probably familiar with the basic requirements, which tend to be common to all of these types of Web applications. You need a database, a Web server and scripting support. XOOPS uses MySQL, Apache and PHP, respectively. You also need some experience with basic Apache, PHP and MySQL server settings.

The first thing to do is to get the latest release of XOOPS from the XOOPS Web site (www.xoops.org). You should be able to find

a tarball or a zip file containing the XOOPS GPL source code. Select the proper mirror, and download the file. Here's a sample download command (you may need to change the name of the file if XOOPS is updated to a newer version):

```
$ wget http://easynews.dl.sourceforge.net/sourceforge/
↳ xoops/xoops-2.0.13.2.tar.gz
```

Most commercial host providers that use Linux for their servers support Apache, PHP and MySQL. So, if you are using a commercial host provider, you won't need to install any of these packages. You will simply need to upload the XOOPS source code to your host server. Depending on the host provider, you may have to unpack all the files first and upload the directory tree, or you may be able to upload the packed XOOPS file and unpack it on the host server.

If you are serving up XOOPS on your own Web server, you have to install the Apache Web Server with PHP support and the MySQL Database Server. You can find Apache packages for virtually every distribution of GNU/Linux, although you may have to install PHP support separately, depending on which distribution you are using.

Now, the first thing to do to get started with XOOPS is to place the XOOPS tarball's content in a folder that Apache can access. The default document directory for Apache often varies from distribution to distribution, but for our example installation, we use the directory `/var/www` (although in our case we will be modifying the default document root for Apache once we have the XOOPS files installed).

Move the tarball and untar it:

```
$ mv xoops-2.0.13.2.tar.gz /var/www/
$ tar -zxvf xoops-2.0.13.2.tar.gz
```

You should notice that you now have a new directory `/var/www/html/`.

In our installation, Apache stores its configuration files in the directory `/etc/apache`. You may find the Apache configuration files in some other directory (such as `etc/httpd/`) depending on your distribution. The easiest way to use XOOPS is to modify the Apache configuration file (usually



ASA COMPUTERS

www.asacomputers.com

1-800-REAL-PCS



Figure 1. The test file shows that PHP is working.

called httpd.conf). Modify the DocumentRoot variable to tell Apache where you put the XOOFS source code. In our case, you want to set the DocumentRoot to /var/www/html/:

```
DocumentRoot /var/www/html/
```

We'll assume you have PHP and MySQL installed and ready to go. If you are running XOOFS on your own server, it is possible but unlikely that you will need to modify the default settings for PHP in order to make XOOFS work properly. PHP uses the file php.ini for its configuration.

You normally use PHP as an Apache module. If your installation hasn't already done so, you can add the following lines to your Apache configuration to get PHP working with Apache:

```
LoadModule php4_module libexec/libphp4.so
AddModule mod_php4.c
AddType application/x-httpd-php .php .phtml
```

You can set up Apache with PHP version 5, but if you do, you might run into some problems with modules and code structures. So, if you don't absolutely need to use PHP 5, use the version suggested by the XOOFS team: PHP 4.0.5 or another minor version of PHP 4.

If you decide to use PHP 5, add the following lines to the httpd.conf file instead of the ones above:

```
LoadModule php5_module libexec/libphp5.so
AddModule mod_php5.c
AddType application/x-httpd-php-source .phps
```

Here is how to test Apache and PHP to see if you have both configured and working correctly. Take a text editor and create a file called /var/www/html/info.php (make sure it has the proper ownership and privileges so that Apache can read and use the file). Put the following lines in the file:

```
<?php
phpinfo();
?>
```

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—\$7,649

- 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.

You may have to modify the privileges. Here's a sample command to do that:

```
$ chmod +xr /var/www/html/info.php
```

If you have Apache running on your own server, point your browser to `http://127.0.0.1/info.php`. If you are using a commercial host, point your browser to `http://yourdomain/info.php`. If everything is working, your browser should show something like what you see in Figure 1.

If you're using a commercial host, the MySQL Database Server will already be up and running.

If you are using your own server, you may need to configure the MySQL Database Server. The configuration file for the MySQL server is called `my.cnf`. You will find it with your distribution, or in the source code. MySQL commonly includes three sample files named `my-large.cnf`, `my-medium.cnf` and `my-small.cnf`. Choose the one that seems right for your needs, and copy it to `my.cnf` as your starting configuration. XOOPS doesn't need any special database settings, so you can use the `my-small.cnf` file as the template for your configuration. Once you are satisfied with the configuration, start up the MySQL server.

XOOPS Installation

Now, we'll assume you have MySQL and Apache/PHP running, so it is time to start with the installation of XOOPS.

Point your Web browser to `http://127.0.0.1/install/index.php` if you are using your own server. Point it to `http://yourdomain/install/index.php` if you

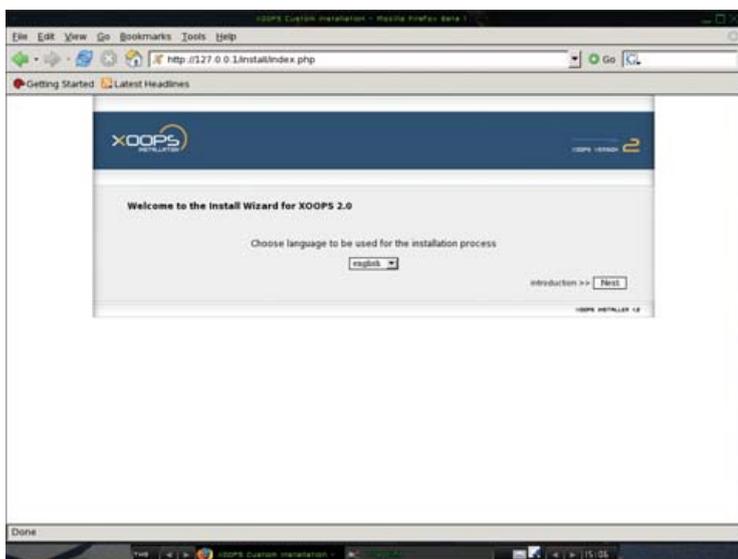


Figure 2. The First XOOPS Install Screen

are using a commercial host.

This should take you to the install screen. See Figure 2 for an example.

Now, click on the Next button to see a brief introduction, as shown in Figure 3.

Click on the Next button once again. At this point, it is quite possible you will experience a problem with the permissions (see Figure 4 for an example).

Here's how to solve these permission problems:

```
$ chmod 777 uploads/  
$ chmod 777 cache/
```

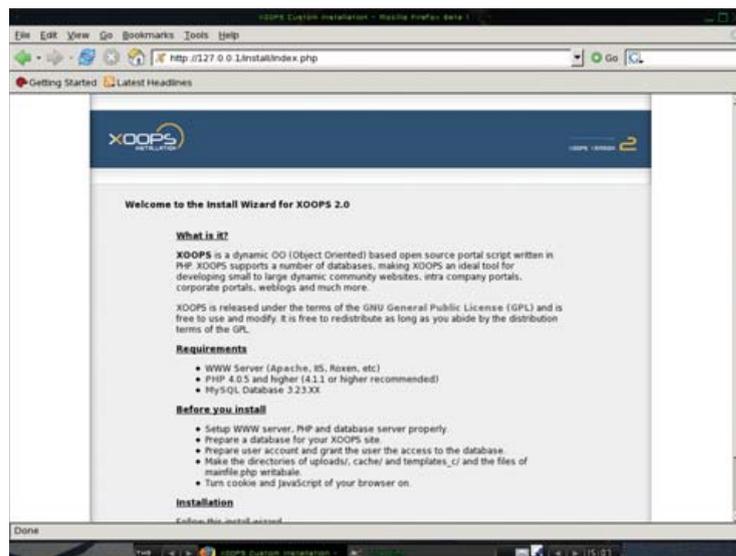


Figure 3. XOOPS Introduction Screen

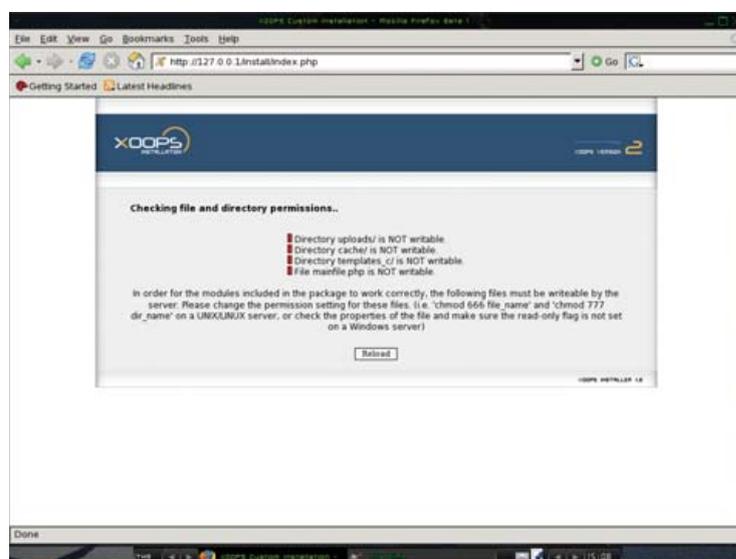


Figure 4. XOOPS reports permission problems during installation.

```
$ chmod 777 templates_c/  
$ chmod 666 mainfile.php
```

Once you have set the correct permissions, press the Reload button, and you should see that there are no longer permission problems (Figure 5).

Press the Next button again. Before we explain how to fill in the next page, you need to create the XOOPS MySQL database and a user name and password for that database. Open a console or an X terminal, and run the following commands:

```
$ mysqladmin -u mysql -p create xoops4  
$ mysql -u mysql -p
```

In this case, I selected `xoops4` for the database name, and I used the default `mysql` as the MySQL user. You can use these names if you want to, but make sure to remember them. You'll need to fill in this information in

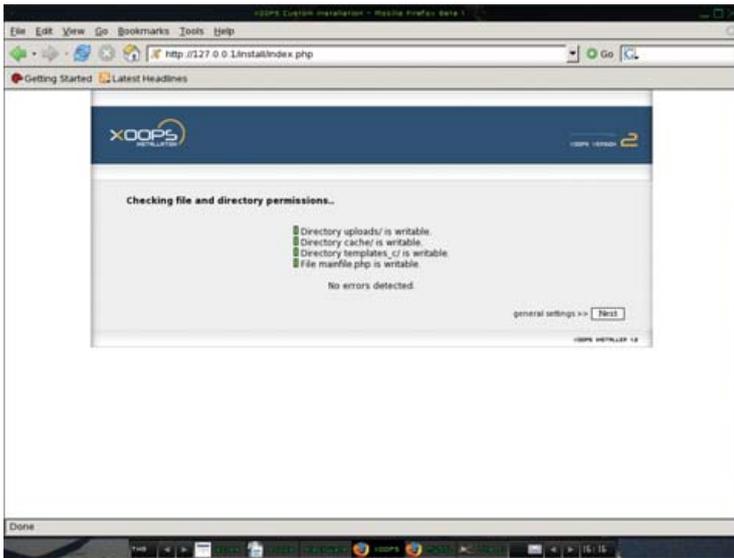


Figure 5. No more permission problems for XOOPS.

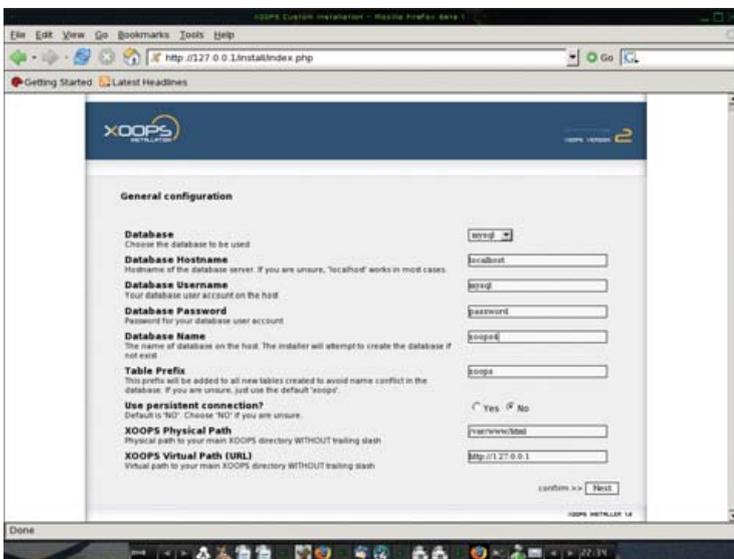


Figure 6. Tell XOOPS about the MySQL database, the database user and password as well as other information.

a Web page shortly. The second command above logs you in to MySQL and presents you with a MySQL command prompt. Once you see the prompt, type the following command to give the mysql user privileges to the xoops4 database, and set the password:

```
mysql> GRANT ALL PRIVILEGES ON xoops4.* TO mysql@localhost
IDENTIFIED BY 'password';
```

Now that you've made a database called xoops4, you must replace its password in the command shown as 'password'.

Now, back to the XOOPS Web configuration page. Fill in the correct information. Figure 6 shows the information for our sample installation. Normally, you would fill in the last field with the domain name for your Web server instead of using 127.0.0.1.



"Company With A Vision"

1-877-25-SERVER : 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 for all your hardware needs.

GENSTOR SYSTEMS, INC.

780 Montague Exp. # 604, San Jose, CA 95131
 Phone: 1-877-25 SERVER or 1-408-383-0120
 Email: sales@genstor.com Fax: 1-408-383-0121

Prices, Products and Availability subject to change without notice.

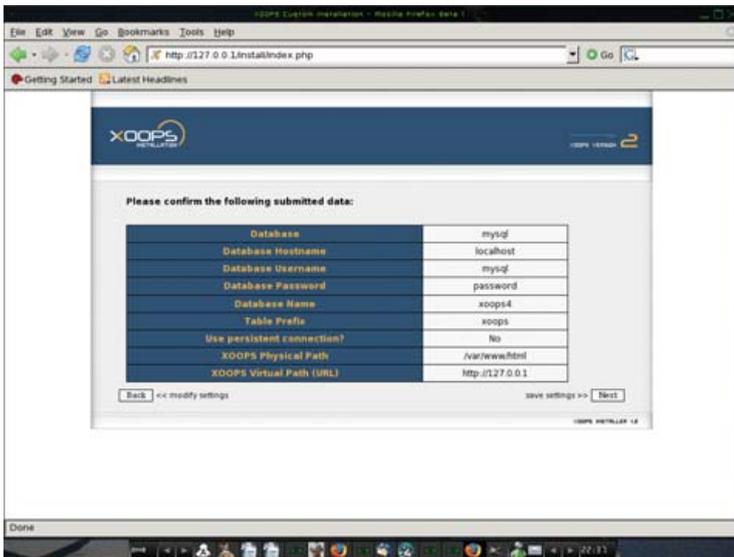


Figure 7. XOOps Information Confirmation Screen

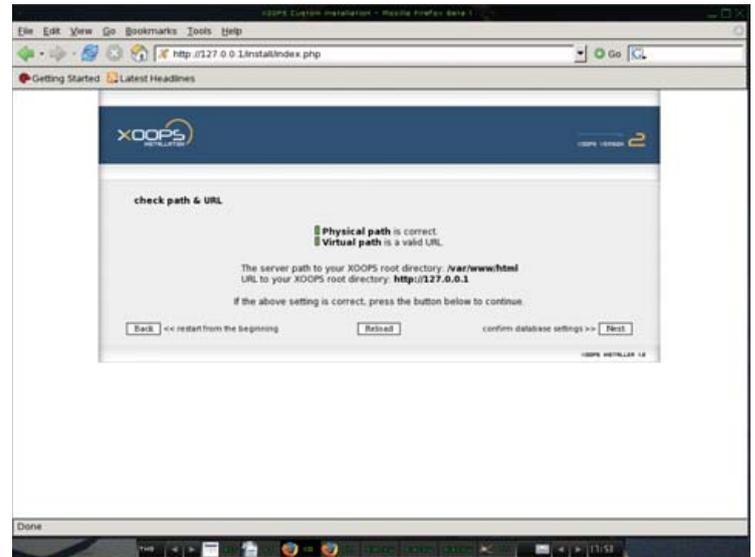


Figure 9. All the paths check out.

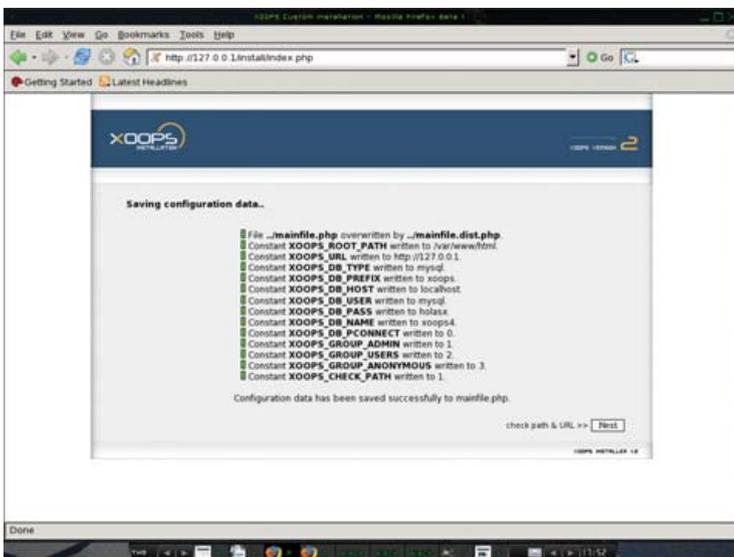


Figure 8. XOOps tells you the configuration data it saves.



Figure 10. Database settings check out.

After you click the Next button, you should see a screen that confirms all the information you entered (Figure 7). If everything looks fine, click Next. If something went wrong, go back and modify the information to correct it.

When you press Next after the confirmation screen, you should see a new screen with the title Saving configuration data (Figure 8). Press the Next button again. XOOps will check to see if the physical and virtual path to the XOOps files are correct (Figure 9).

Click the Next button again. You should see a Web page that confirms the database settings (Figure 10).

Next, XOOps checks to see if it can connect to the database and modify the database (Figure 11).

Click the Next button once again. Now the XOOps installer creates all of the MySQL tables it needs (Figure 12).

Now you need to fill the text boxes with information about the administrator of the XOOps site (Figure 13). XOOps will ask you about the

Admin name, Email and the password for the Admin account.

Click Next, and XOOps uses this information to insert some starting data in various XOOps tables (Figure 14).

Click the Next button once again and that should get your XOOps site up and running. You should see a final confirmation screen like the one shown in Figure 15.

Press the HERE link to get to the login screen for your site. This is actually what your site looks like when you first start XOOps. It is obviously bare bones, but that is because you haven't activated any of the XOOps features yet.

Log in using the administrator user name and password you specified earlier.

XOOps Administration

Go to the Administrator Menu. You will see a number of menu icons on the left. These options configure various features of XOOps.

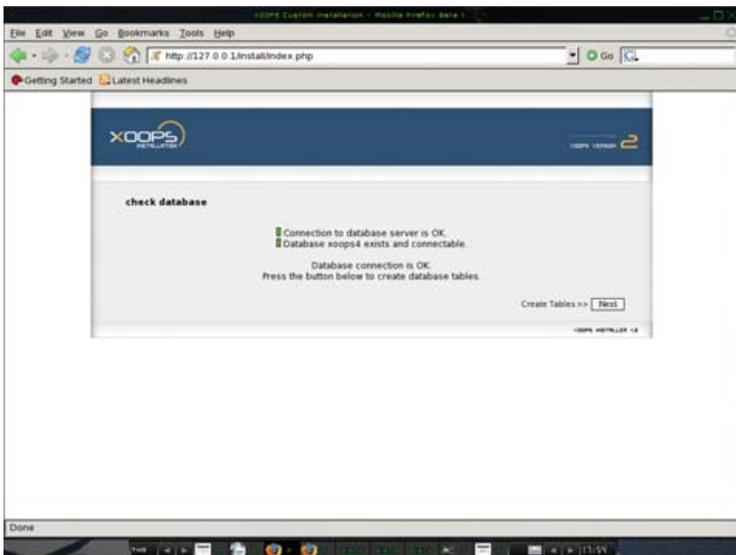


Figure 11. XOOPS can connect to and use the database.

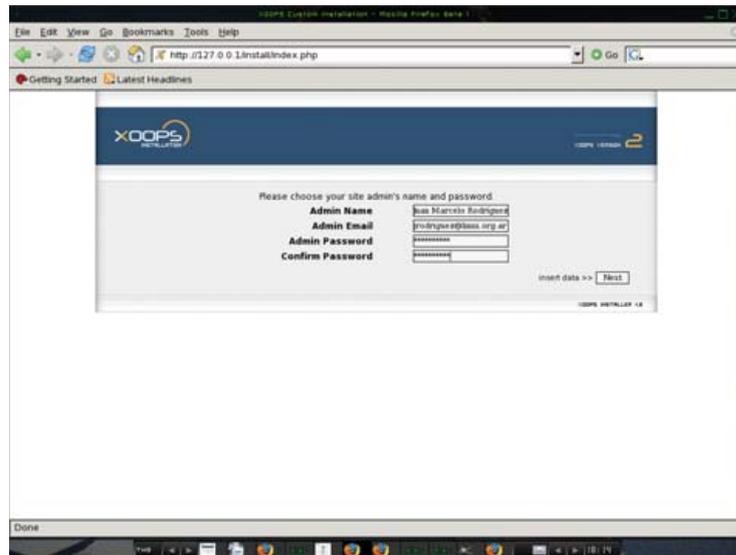


Figure 13. Specify the XOOPS Administrator e-mail, user name and password.

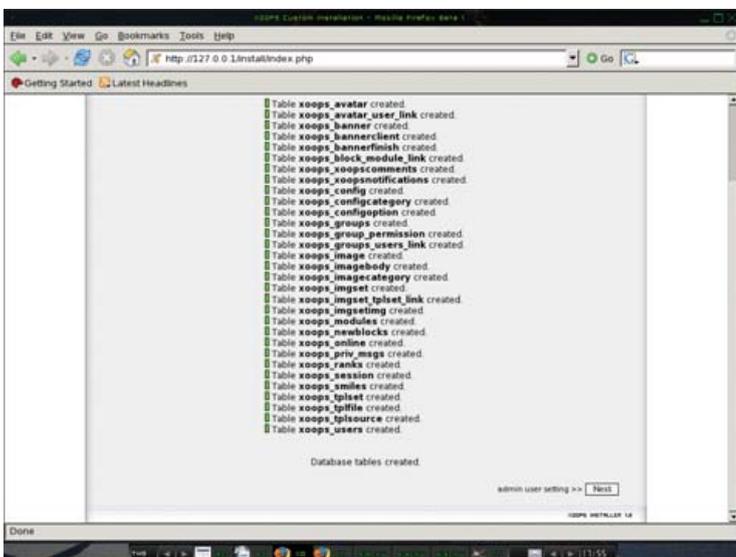


Figure 12. XOOPS creates all of its own database tables.

Due to the copious amount of files and information that are currently available for this application, it would be almost impossible to describe each module and plugin in-depth in this same article. However, according to my humble opinion, one of the most significant features—and also the first thing you will need to set up—is the option Preferences. On that menu, you can find the following submenus: General Settings, User Info Settings, Meta Tags and Footer, Word Censoring Options, Search Options and Mail Setup.

I strongly recommend that you start working with General Settings first by clicking on the Edit tag. Now you should see another name displayed when you click on System Admin and its modules.

From this menu, you can activate and deactivate your modules, or just add new ones. You can download those new modules from the XOOPS site.

Go to the Administrator Menu, then to the Modules section and see how it looks (Figure 17). Here is where you can install, uninstall, activate and deactivate XOOPS modules.

The next thing you should do is set up some basic information about your site. Go to the Administrator Menu, choose Preferences and then General Settings. Specify your Site name, Site Slogan, Theme, Admin mail address, Time Zone and so on.

Once you have the modules installed and running, you can make the

LAYER 42™

- Redundant UPS and generator
- Nationwide network
- Free tech support

2U 256kbps ~80GB \$60/mo.	4U or Mid-tower 256kbps ~80GB \$80/mo.
1/4 Rack 512kbps (14U) ~165GB \$200/mo.	1/2 Rack 1mbps (28U) ~330GB \$350/mo.

www.layer42.net

All prices include 100Mbps port, Firewall,
24x7 Monitoring and DNS hosting

408-450-5740
2312 Walsh Ave., Santa Clara, CA 95051

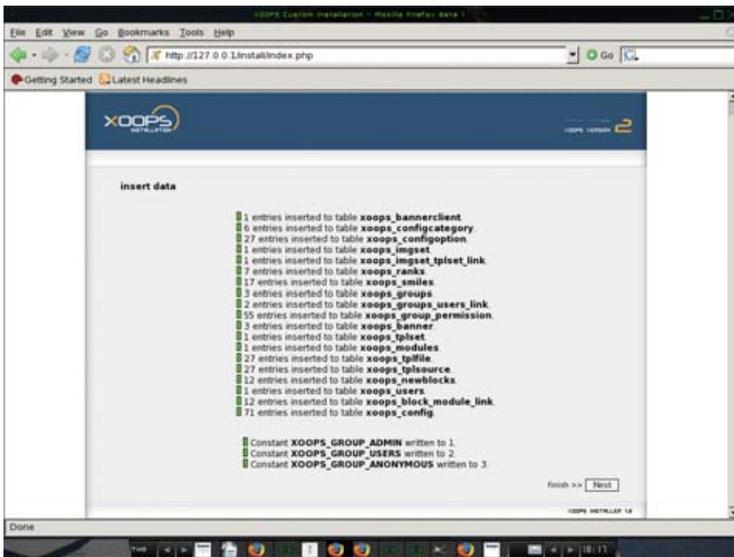


Figure 14. XOOps tells you about initializing data in various tables.

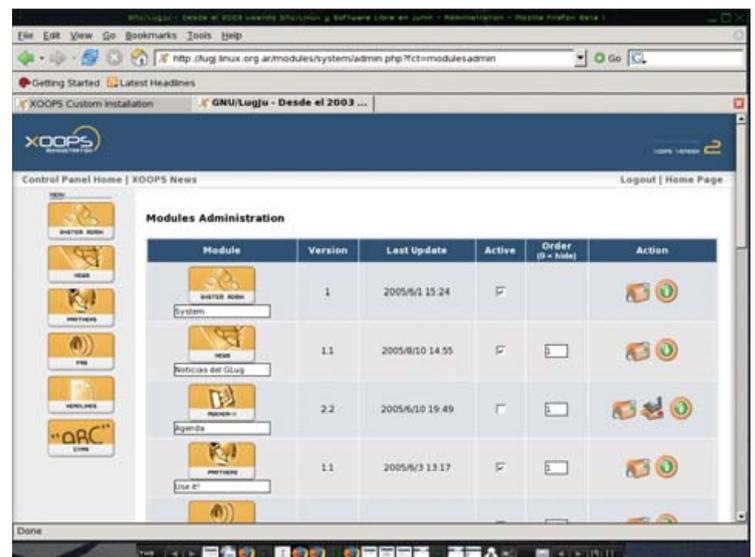


Figure 17. Manage Your XOOps Modules



Figure 15. Final XOOps Installation Screen

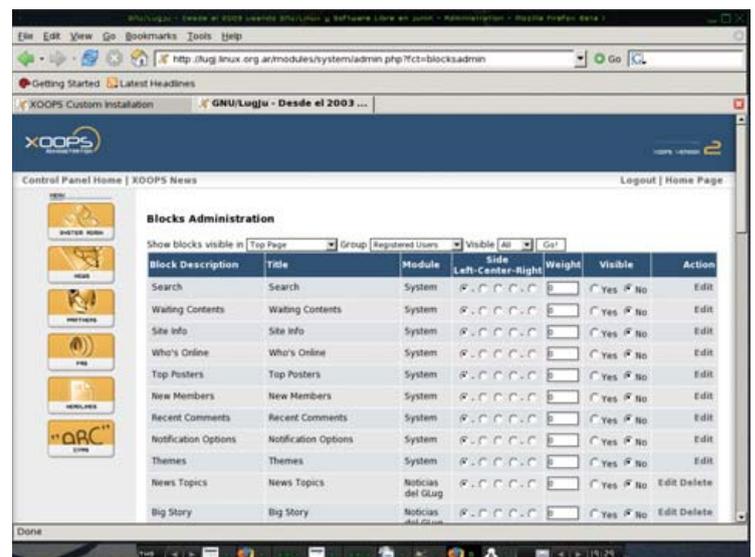


Figure 18. Configure which blocks you want visible and where they should appear.

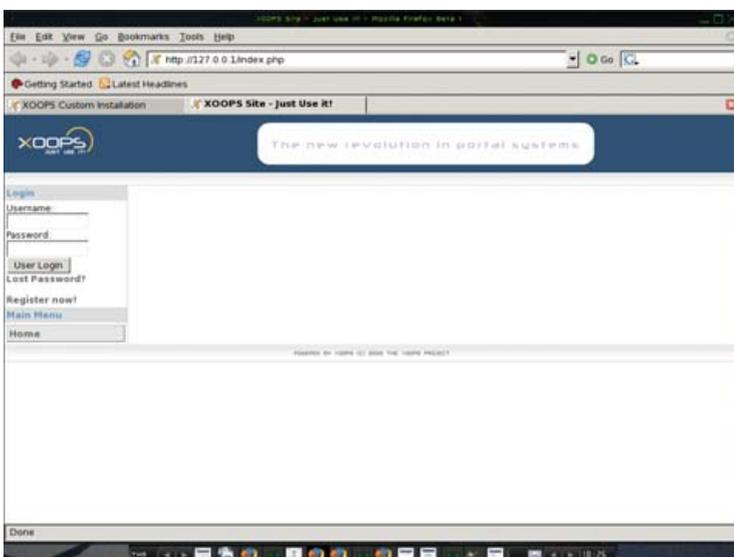


Figure 16. Login Screen for Your XOOps Site

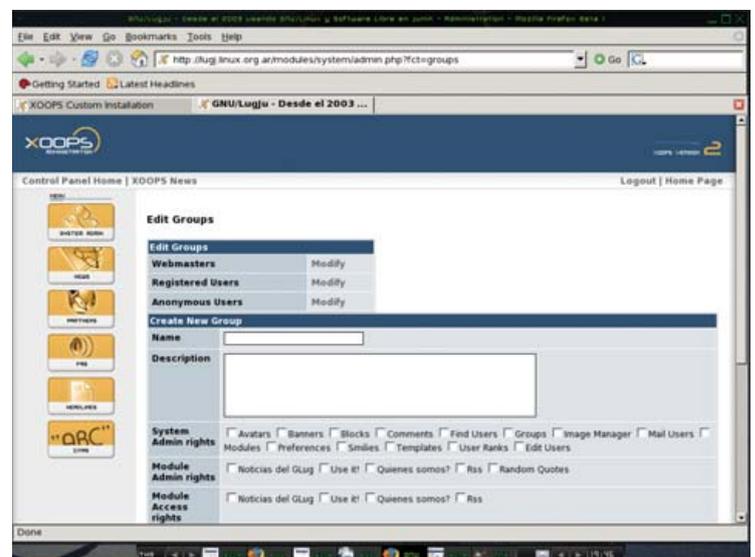


Figure 19. Create and manage user groups for your site.

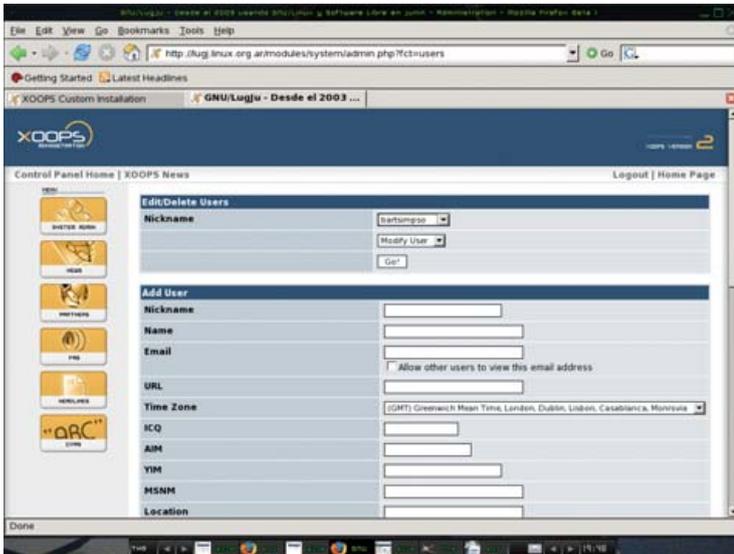


Figure 20. XOOPS User Management

features they provide show up on your XOOPS home page. Go to the Blocks Administration section. You can add, modify or activate whatever modules you want and the features visible in various locations (the left column, middle, right column and so on). You can specify who will be able to see or use the various features. In most cases, you will want to let everyone view the blocks you choose to display. See Figure 18 for a sample blocks configuration screen.

There are cases where you might want to restrict who gets to see or use blocks. You do this by defining groups of users and assigning people who register for your site to the various groups. You can set, modify and add groups through the section Groups in the Administration Menu (Figure 19).

XOOPS allows users to register for your site. In this case, XOOPS adds the users to your database. To add users yourself, or modify existing users, check out the option Edit Users (Figure 20).

Adding More Modules and Themes

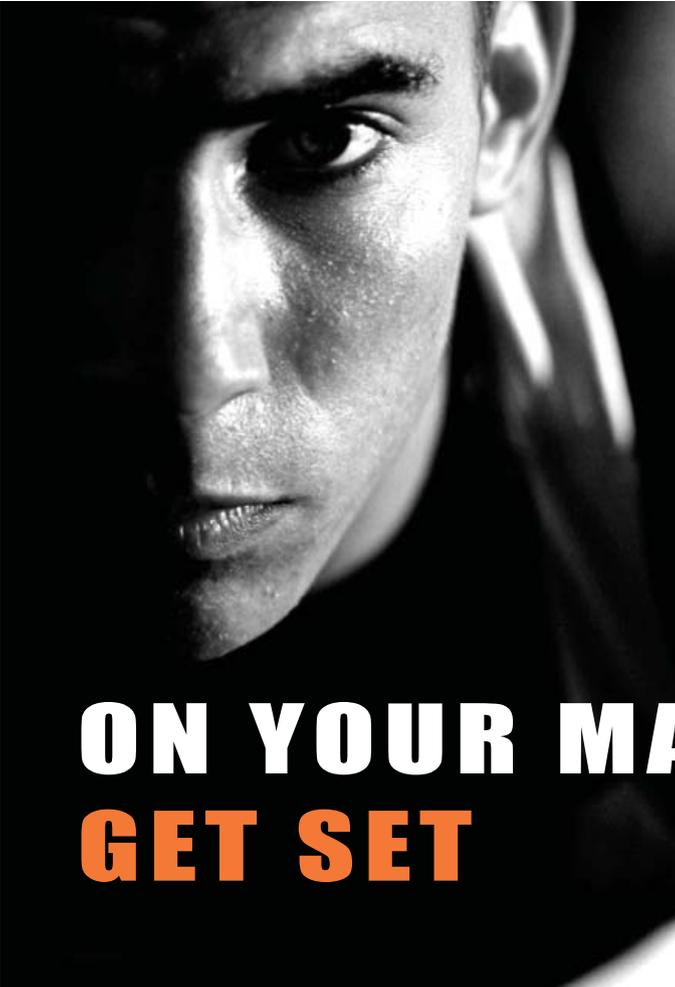
XOOPS keeps modules in the directory `/var/www/html/modules/` in our case. Obviously, if you use a different document root, you'll have to modify this path to suit your installation. If you want to add a new publicly available module to XOOPS, download it from the XOOPS site or from another resource, then unpack it in the directory for the modules (in our case, unpack it in `/var/www/html/modules/`). Go to Administrator Menu, then to modules, and you will have the option to install and activate this new module.

You can also download custom themes for XOOPS. In our case, you would unpack a custom theme to the directory `/var/www/html/themes/`. You can set this new theme as the default theme in the General Preferences section.

There are many other things you can do to customize your installation of XOOPS. Work and play with the various administrator tools—as you usually will discover the best features by experimentation.

Now you have a powerful Web site, based on a terrific object-oriented language and running upon a strong database and Web server and, of course, the most robust operating system available, GNU/Linux. ■

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.



ISPCON

The Definitive Event for Wired & Wireless ISPs

May 16-18, 2006 • Baltimore Convention Center

www.ispcon.com

ISPCON is the starting block for the next phase of success in your company. Make those all-important connections with your peers and predecessors, and come away with the tools and technology that will build your business.

Register today with **Customer Code LNXS6**
to get your **FREE EXHIBIT HALL PASS!**

**ON YOUR MARK
GET SET**

SUCCEED

Remote Temperature Monitoring with Linux

Use a small footprint Linux with some cheap hardware to create a remote temperature monitor. STEVEN M. LAPINSKAS

I started this project to record and access temperature readings remotely after I had a conversation with a friend who is in the HVAC business. His job is to make sure the climate indoors is comfortable—not too hot and not too cold, depending on the weather outside. He finds many new installations have startup bugs that must be worked out, because no two installations are exactly the same. The end of the job is the most stressful. A customer calls after he is done and lets him know something is wrong. The customer isn't happy, but he doesn't know where to start looking for the problem because there isn't any good objective information about what is going wrong with the installation.

We agreed that it would help to be able to record the outside temperature and log the readings electronically. This would be one way of improving the troubleshooting process. I then started the search to purchase an off-the-shelf recorder that was inexpensive, easy to install and simple to use. While looking, I found a wide range of commercial products and kits. Some are standalone and some use a PC for displaying and recording temperature data. Each of our three requirements was equally important and I found that most of the products were too expensive for our budget. Ease of installation was typically another problem. Some devices had complicated wiring or the requirement that they be placed where temperature was measured. Many people consider thermostats to be just clutter on a wall, so it wasn't going to be easy to convince them to have another box to record temperature.

The Linux Solution

Finally, to solve the problems of cost and installation, I looked at the possibility of building a system from components. Following the Linux idea of assembling and integrating tools to get a task done, I looked at using a digital multimeter, a PC and software to make them work together. The digital multimeter with an RS-232 serial port interface would measure temperature using a sensor. The PC would collect the data from the multimeter and process it for display.

I was aware of open-source utilities for the multimeter serial port interface using Linux and had purchased a multimeter earlier for general troubleshooting. We had a retired PC available, so all the components were on hand to build a prototype temperature recording system.

The digital multimeter came with a serial port cable and DOS software. I didn't use the supplied DOS program. There was no way to modify the program to allow temperature measurements with the sensor. Instead, I used QuickBasic to write new software from the ground up. I had the necessary details about the serial port

interface for the multimeter, and QuickBasic had all the features I needed. I got a prototype communication program to work, but I ran into memory management issues with DOS and QuickBasic as the application grew in size, especially when I started dealing with the need to display and record data.

It seemed like a big step backward to struggle with memory management at this point. I knew Linux would provide an environment where I wouldn't need to be concerned with memory management, so I looked for a distribution to use as a replacement for DOS.

I found that the most popular Linux distributions weren't appropriate for this application. Even a minimal installation of these distributions would exceed the capacity of the retired PC. The distribution I found that overcame these restrictions was University Linux from Paul Muller. It has small memory and disk requirements. I was able to run it on the retired PC using less than 20MB of DOS formatted hard disk space and 24MB of RAM. Best of all, the distribution is tolerant to power failures. If the power goes out, the PC reboots without causing file corruption problems that need manual help. This saves money and reduces complexity, because I didn't need a UPS to keep the system running during power failures.

Once I configured everything on the PC, there was no need for a keyboard or monitor. I could use a Windows PC and Telnet, along with an Ethernet connection to communicate with the system PC for development and testing. I prefer to write and test incrementally, so I chose Perl for the language for this project. University Linux comes with Perl version 5.003. I couldn't use Perl modules, the application size was too small, so this was a minor inconvenience. University Linux also includes Acme Labs tftpd server. This allowed me to set up the system to use a Web browser for viewing temperature measurements.

Testing Hardware

I used a Tandy Catalog No. 22-805 digital multimeter that comes with an operating manual, DOS software, wire test leads and serial cable with nine pin connectors. According to the manual, the communication settings are 600 baud, seven data bits, two stop bits and no parity. Important information was left out of the operating manual, but I found what I needed on the Web. The DTR and RTS lines need special attention. The DTR line has to be set low and the RTS line set high for the meter to communicate through the serial port. It is impossible to get data from the meter without the two lines set this way.

I could use only stty for serial communication with this distribution and couldn't explicitly control the DTR and RTS lines in the script. This meant I needed a hardware hack to make things work.

I found that DTR and RTS change from a low to high state when I call stty in the script. This works out okay for DTR, but RTS has to remain low. I realized that the second serial port on the PC has RTS low as it isn't being used. If I connected the multimeter serial interface cable RTS to the RTS pin of the second serial port, the multimeter would be faked into seeing the correct line setting. I simply removed the RTS line from the multimeter and connected it to the second serial port.

With that problem solved, I powered up the multimeter and put together a short test script (serialtest.pl), as follows:

```
#!/usr/bin/perl
#
# serialtest.pl
#
# Script for reading Tandy Model 22-805 meter
# through serial port.

$port = "/dev/ttyS1"; # set to COM1

system ("stty 600 cs7 cstopb clocal -ixon -echo < $port");

open (SERIALPORT, ">+$port") or die "can't open $port. ";

print SERIALPORT ("\n"); # take a reading

$R = <SERIALPORT>; # read returned string

print "$R" ;

close (SERIALPORT); # close port

exit 0;
```

If the script ran successfully, I would get a string with the same reading shown in the multimeter LCD. I set the multimeter to the resistance measurement range and ran the script. The result was:

```
OH 0.L MOhm
```

A good start! The hardware hack worked. Now it was on to measure a temperature sensor with the multimeter.

I chose an NTC (negative temperature coefficient) thermistor for the temperature sensor. Despite the fearsome sounding name, this is just a small two-wire electronic component that changes electrical resistance with temperature. With a multimeter, the resistance measurement provides information to tell temperature. The thermistor is impossible to wire backward, because it isn't voltage-polarity (+ or -) sensitive. This means one less thing for the technician installing it in the field to worry about.

The thermistor isn't fragile, but the leads to the

body can be broken with excessive tugging or bending. I used a two-position terminal block to solve this problem and make the connection to the wiring simple. I placed one thermistor lead and a wire under a screw terminal and then tightened the screw to make a solid mechanical and electrical contact.

With the thermistor connected to the ends of the test leads and the test leads plugged in to the multimeter, I powered up and ran the test script again. The result was a resistance reading:

```
OH 34.23kOhm
```

The numeric portion of the reading is 34.23 with a k after it. The k is an abbreviation for kilo or 1,000. Because the multimeter LCD doesn't have enough characters to display large numbers, it uses a multiplier. In this case 34.23k is 34,230 Ohms.

I found that this reading was very close to 0°C by referencing a table of resistance-to-temperature values supplied by the manufacturer. It matched the temperature reading of another thermometer with a sensor in the general area, so I was confident that this assembled system would work and provide accurate readings.

Now it was time to create a script to use the data and display the temperature value.

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

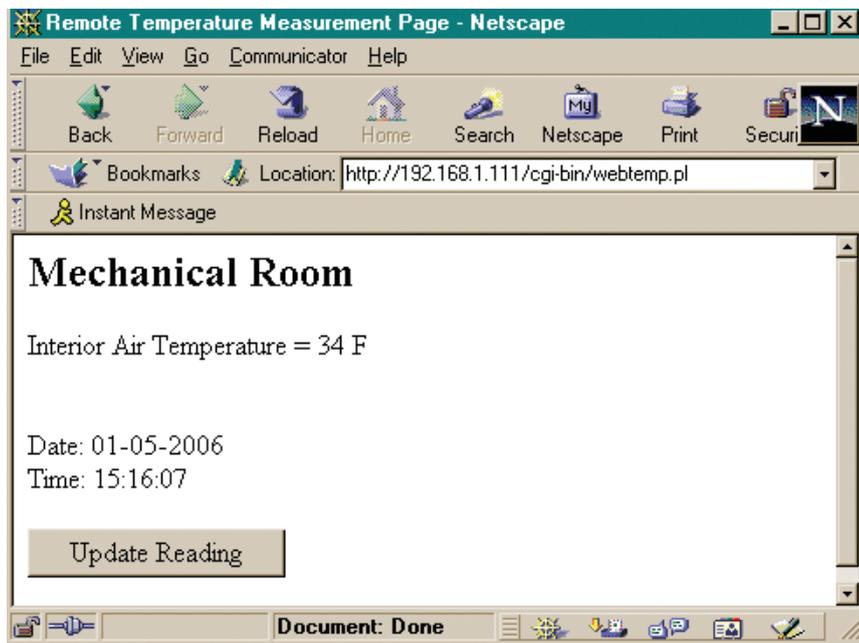


Figure 1. The temperature monitor now has a Web interface.

Software Choices

Two choices were available to perform resistance-to-temperature conversion in the script. I could use a lookup table with pairs of resistance-to-temperature values in an array. The sheer number of elements in this array would be a drawback to this approach. A span from -40°C to $+40^{\circ}\text{C}$ requires 81 (don't forget 0°C) pairs of values. There was no easy way to manipulate a text file available from the thermistor manufacturer, and entering the values by hand would take time and be prone to errors.

Instead, I used what's called the Steinhart-Hart equation (see sidebar). The equation was developed in the late 1960s to help process ocean temperature data collected with thermistors and provides direct conversion of resistance to temperature. A spreadsheet utility found on the Web helped with calculating coefficients unique to each family of thermistors and was used in the equation.

Display Data

Once the script calculates temperature from a multi-meter reading, it needs to be displayed or stored. With this in mind, I extended the test script to convert and display temperature, and show the time and resistance reading. University Linux uses the 2.0 kernel, and root user login by Telnet is allowed. When ordinary users attempt to run the grabtemp.pl script, an error is displayed because of the file permissions used for the serial port, /dev/ttyS1. I fixed this by changing permissions with:

```
chmod a+x /dev/ttyS1
```

Now, ordinary users could log in and run the script to check temperature. They wouldn't need root access.

Here is the output from the resulting showtemp.pl script:

```
/perlserial: perl -w showtemp.pl
01-05-2006 14:43 34 F 1.3 C 30.52 k Ohms
```

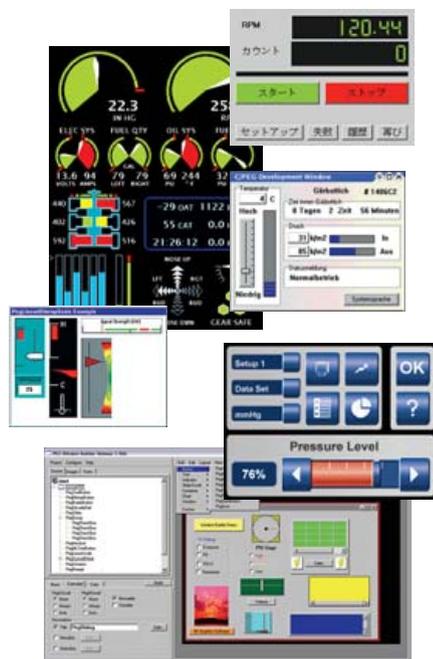
Here you can see the date, time, temperatures in degrees F and degrees C, along with the actual resistance reading. I checked the temperature where the sensor was located and found that the reading was accurate, so the conversion formula part of the script worked.

Not too many computer users are comfortable with using a command-line program interface. Web browsers with a point-and-click interface are a lot less intimidating. So, I extended the script once again to allow users to operate the system with a Web browser.

With the tthttpd server configured and running, it was just a matter of directing the output from the script to build a Web page for display. This was fairly straightforward as the following code shows:

PEG[®] the leading GUI for Embedded Systems

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



Prototype Today!
FREE EVALUATOR PACKAGE



GRAPHICS SOFTWARE FOR EMBEDDED SYSTEMS

WWW.SWELLSOFTWARE.COM 810-982-5955

World's First 1U 64GB Server RS-1564 Open Series



Why you should choose OSS ?

- **Customer Service:** OSS has an industry leading 98% customer retention rate. Our success relies on your satisfaction.
- **Rapid Delivery:** OSS ships most orders from 1-1000 systems in 3-6 days. OSS designs and manufactures all systems in our state of the art multi-million dollar, fully-automated 32,000 ft² production line.

Find Out More »

1-866-664-STOR
sales@OSShpc.com

OpenSourceStorage.com

OSS is not just another vendor...
We're your partner in business.

Open
Source Storage®

**AMD**
Solution Provider
GOLD

©2003 Advanced Micro Devices, Inc

©2002-2006 Open Source Storage, Inc. All Rights Reserved. AMD, AMD Opteron, combinations thereof, are trademarks of Advanced Micro Devices, Inc.

Digital Edition Now Available!

Read it first

Get the latest issue before it hits the newsstand

Keyword searchable

Find a topic or name in seconds

Paperless archives

Download to your computer for convenient offline reading

Same great magazine

Read each issue in high-quality PDF

Try a Sample Issue!

www.linuxjournal.com/digital



```
print "content-type: text/html \n\n";
print "<HTML><BODY><P>";
print "<HEAD><title>Remote Temperature Measurement Page</title></HEAD>";
print "<H2>Mechanical Room</H2> ";
print '<form action="webtemp.pl" method=post> <P> <P>';
print "Interior Air Temperature = $out_tempF<BR>";
print "<BR>";
print "<BR>";
print "Date: $out_date <BR>";
print "Time: $out_time <BR>";
print "<BR>";
print '<input type=submit value="Update Reading">';
print "</form>";
print "</BODY></HTML>";
```

Running the webtemp.pl script from /cgi-bin gives the user a display like the example shown in Figure 1.

This example shows the temperature in the room as well as the time and the date of the reading. You can press the Update Reading button to rerun the script and display another temperature value.

It is easy to write an extension to the script to log temperature over time. I put a line in the rc (boot)

script that launches a data logging script, which then runs continuously in the background. I found that I could use measurement intervals of 5–10 minutes, because changes in air temperature are slow indoors in an air-conditioned space.

You can access the temperature log through the command line by using Telnet. Because the format was space-delimited, the date file was used with Microsoft Excel to plot graphs and view trends. You can see a sample output in Figure 2.

Security Concerns

The overall objective was to create a reliable and easy-to-use electronic means to display and record temperature data. When you actually deploy the system, the location of the system and the network connection can vary widely. Depending on circumstances, you have to evaluate the security

```
192.168.1.111:23 - KevTerm
Connection Edit Options Help
/perlserial: tail -c 225 thermistor.log
01-04-2006_19:51 28.1 F -2.1 C 36.35 k Ohms
01-04-2006_19:56 29.0 F -1.6 C 35.41 k Ohms
01-04-2006_20:01 28.3 F -2.0 C 36.15 k Ohms
01-04-2006_20:06 28.4 F -2.0 C 36.09 k Ohms
01-04-2006_20:11 28.0 F -2.2 C 36.44 k Ohms
/perlserial: █
Ready Ln 7, Col 14 69x14
```

Figure 2. A sample set of values kept in the temperature monitoring log as seen through Telnet.

THERMISTORS AND STEINHART-HART EQUATION

The plot of resistance to temperature on a graph for a thermistor looks a lot like the curve of a ski jump, and each family of thermistors has their own unique curve. So, simple $y = mx + b$ algebra won't help to convert resistance to temperature. The equation of the curve can be described by a polynomial. The Steinhart-Hart equation is a trinomial or an equation with three terms. Solving the equation at each resistance measurement point requires three coefficients; a , b and c . Some manufacturers provide these for their thermistors. Others provide only tables of resistance to temperature.

When the coefficients aren't available, a spreadsheet utility is available to help find them using the manufacturer's tables. No algebra is required; simply enter the three values of temperature and resistance into the spreadsheet and the coefficients are calculated automatically. They are usually very small numbers, expressed in scientific notation. But the calculated coefficients can be cut and pasted from the spreadsheet into the Perl script for use, reducing errors from typing

concerns for each installation. You may have to implement some workarounds to address the security concerns. For example, you can log temperature readings in the form of text or HTML pages by a script running in the background and not by a script in the cgi directory, which isolates the logging process from Web access. Alternately, you can gather data from this server using another secure server through FTP or HTTP. This would add another layer to prevent direct access by the outside world, but still make the information available.

Future Extension

Digital multimeters are general-purpose electronic measurement tools. Although I used a thermistor for temperature measurement in this application, you can use other sensors that have resistance, voltage or electrical current as outputs. Some other conditions to measure include flow, pressure, weight, light level and humidity.

You don't need more multimeters to measure more than one temperature. You can connect a single multimeter to a switching device. You then would create a script to operate the switching device, which allows you to select one temperature sensor at a time.

Conclusion

This example shows how the tool concept behind Linux works for solving applications where cost and flexibility requirements are important. The wide variety of distributions available compared with other operating systems meant developing a system with all the features needed was practical. Additionally, you can add features using Perl and the development environment provided by the University Linux distribution.

The system can be duplicated for less than \$100 US. The multimeter, thermistor and wiring accessories are available from numerous electronics retailers. Many retailers have Web sites, so it's easy to compare features, specs and pricing before ordering. Purchasing a used digital multimeter should be done with caution, as there is no easy way to tell whether accuracy of the instrument has been affected by the previous use.

Resources for this article:

www.linuxjournal.com/article/8833 ■

Steven M. Lapinskas has a professional background that includes the areas of software quality assurance, mechanical design and project management. Some of his free time is spent experimenting to interface Linux with the real world outside the computer.

www.faircom.com/go/?track

Your packages get
**tracked &
delivered**
on time by FedEx.

FairCom database
technology makes
it possible.

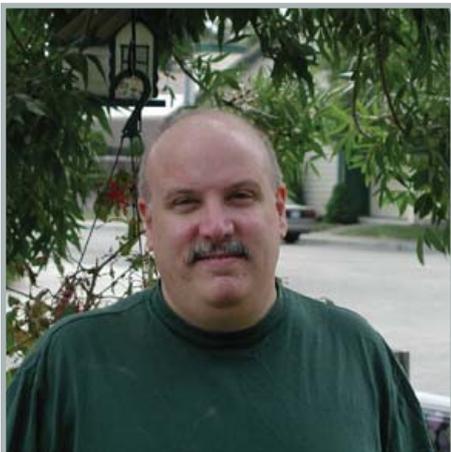


FairCom®

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

The Spirit of Open Source

How dare Linus Torvalds recommend that people use what's best?



Nick Petreley, Editor in Chief

Dang. I had intended to rant about wikis this month, but Dave Taylor covered the topic thoroughly and did a much better job than I would have. See his article “Why I Don’t Like [edit: Get] Wikis” in this issue.

There’s plenty of other things to rant about, however. There’s the schizophrenic, religious and hypocritical zealotry of free/open-source advocates that often gets more ink than the sane attitudes that are more prevalent in the development community itself.

Take the irrational fear of Java and its gatekeeper, Sun, as an example. Do you realize there are people who still insist that the only acceptable version of Java is a clean-room open-source implementation that (they presume) cannot be controlled by Sun? Did you also know that, according to Evans data, the vast majority of Linux developers uses Java-based Eclipse as their favorite integrated development environment (IDE)?

I’m using the Java-based Jedit to write this column. I use Jedit because I think it is the

best editor on the planet. Ask me if I’m afraid that Sun will send the Java police after me to collect a license fee. No, ask me what I would do if Sun did that? I’d gladly pay up. Why? I told you. I think Jedit is the best editor on the planet and I want to use it.

Do you know what Linux developers named as their second favorite IDE? KDevelop. That’s right, the KDE-based IDE that depends upon the evil Qt. Sure you can use the GPL version of Qt, which requires you to share your code. But Qt is evil because you have to pay license fees to its creator, Trolltech, if (and only if) you want to sell a closed-source proprietary application based on Qt.

GTK, on the other hand, is good, because you can sell closed-source proprietary applications based on GTK without having to give anything back to the people whose work you exploited in order to make your money. Don’t take my word for it. When I talked to Ximian’s Miguel de Icaza, he named the LGPL license as the reason why people should choose GTK and GNOME over Qt and KDE. And it is the LGPL that allows people to exploit the work of the developers of GTK and GNOME without having to compensate them with money or source code.

In view of this, it is beyond me how GTK and GNOME remain the poster children of open source for so many open-source advocates.

What is the spirit of open source? It is the GNU General Public License. The idea is that if you publish software that integrates someone else’s publicly available work (work licensed under the GPL), you are required to make your additional work available to the public as well.

The Linux kernel is based on the GPL. NVIDIA violates the GPL because it keeps some of its Linux kernel driver code secret. The end result is that you will “taint” the kernel if you use NVIDIA’s closed-source kernel module. Shame on NVIDIA. It isn’t sharing like it’s supposed to.

Fine. I agree with that. But how can you go from there to saying GTK is good because it allows—no, invites—you to do what NVIDIA does? The whole point of the LGPL is to allow you to add something to GTK without having to compensate the GTK developers with either money or source code.

Don’t get me wrong. Personally, I couldn’t care less what motivates people to use Qt, GTK, Java, Python or the practically useless GCJ (GNU Java compiler). What irks me is when someone advocates inferior solutions purely in the name of open source, especially when those so-called open-source solutions so clearly violate the spirit of open source.

If you want a good example of the right attitude, look no further than Linus Torvalds, Linux creator. You don’t have to agree with his methods or his decisions, but I don’t see how anyone can impugn his motives. Here is a man who cares about what’s right and what has practical value.

So what are we to make of the fact that Linus Torvalds criticized GNOME and recommended KDE? Here we have the creator of the Linux kernel criticizing what many see as the poster child of open source and recommending the evil Qt-based KDE. Why would our open-source hero say such a thing? Because in his opinion (an opinion I share), the GNOME design is so bad it should be considered a disease.

You don’t have to agree with him, but it’s plain that his recommendation is based on his opinion of what works best. There’s no sign of misguided zealotry or religion in that recommendation. Use what’s best. What a concept. Linux developers seem to get it. It’s about time the open-source zealots got it too. ■

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



Rackspace – Managed Hosting Backed by Fanatical Support™

Fast servers, secure data centers and maximum bandwidth are all well and good. In fact, we invest a lot of money in them every year. But we believe hosting enterprise class web sites and web applications takes more than technology. It takes Fanatical Support.

Fanatical Support isn't a clever slogan, but the day to day reality our customers experience working with us. It's how we have reimagined customer service to bring unprecedented responsiveness and value to everything we do for our customers. It starts the first time you talk with us. And it never ends.

Contact us to see how Fanatical Support works for you.

1.888.571.8976 or visit www.rackspace.com



Thanks for honoring us with the
2005 Linux Journal Readers' Choice Award for
"Favorite Web-Hosting Service"



The World's Fastest InfiniBand Switch

From a Company You've Trusted for 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.

Working with PathScale's InfiniPath HTX Adapters, the number of hops required to move MPI messages between nodes is reduced, improving latency. The modular design makes them useful for SDR, DDR and future QDR InfiniBand fabrics, greatly extending their useful life. Please send email to fastree@microway.com to request our white paper entitled *Low Latency Modular Switches for InfiniBand*.



▲ 72 Port FasTree™ Configuration

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.

◀ 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