

ASTERISK CALLER ID | AJAX | ATOM | DRUPAL | DOCBOOK | SMART

LINUX JOURNAL

Since 1994: The Original Magazine of the Linux Community

NOVEMBER 2006 | ISSUE 151 | www.linuxjournal.com

WHY AND HOW
PLANETIZEN
MIGRATED
TO DRUPAL



LINUX JOURNAL INTERVIEWS

TIM BRAY

GET CALLER
ID ON YOUR
WEB PAGE
WITH AJAX

A SMART WAY
TO MONITOR
AND REBOOT
SERVICES

PLUS

FILTER PORN ON
MULTIPLE CLIENTS
THE EASY WAY

A CARRIER GRADE
LINUX SUCCESS STORY

RUN SERVICES ON
USER MODE LINUX

AN
SSC
PUBLICATION

USA \$5.00
CAN \$6.50



0 71486 03102 4

Enterprise and High-Performance Computing Under Your Control



Industry Leading 2P and 4P x86 Computing

Innovative server technology with outstanding performance and memory scalability

2-way XtremeServer™ up to 64GB of memory



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

4-way XtremeServer™ up to 128GB of memory



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



Keeping Your Cool in the Data Center while Consolidating & Virtualizing your IT Infrastructure
Go to <http://www.appro.com/whitepaper/White Papers.asp>

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

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



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

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

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



CONTENTS

NOVEMBER 2006
Issue 151



ON THE COVER: PHOTO BY JOHANN WALL

FEATURES

50 INTERVIEW WITH TIM BRAY

Tim Bray releases Atomic energy.

James Gray

54 CALLER ID WITH ASTERISK AND AJAX

Want to do your call screening by Web page?

Mike Diehl

60 MIGRATING TO DRUPAL

What drove Planetizen to migrate to Drupal?

Abhijeet Chavan and Michael Jelks

66 SIMPLE WEB SITES USING DOCBOOK, XML AND CSS

David Lynch

ON THE COVER

- **Why and How Planetizen Migrated to Drupal, p. 60**
- **Linux Journal Interviews Tim Bray, p. 50**
- **Get Caller ID on Your Web Page with Ajax, p. 54**
- **A Smart Way to Monitor and Reboot Services, p. 76**
- **Filter Porn on Multiple Clients the Easy Way, p. 93**
- **A Carrier Grade Linux Success Story, p. 70**
- **Run Services on User-Mode Linux, p. 36**

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.



CoyotePoint
Systems Inc



CONTENTS NOVEMBER 2006

Issue 151

COLUMNS

20 REUVEN M. LERNER'S AT THE FORGE

Beginning Ajax

28 MARCEL GAGNÉ'S COOKING WITH LINUX

The Dynamic Web: for Those Who Like to Watch



32 DAVE TAYLOR'S WORK THE SHELL

Analyzing Log Files Redux

36 MICK BAUER'S PARANOID PENGUIN

Running Network Services under User-Mode Linux, Part I

40 JON "MADDOG" HALL'S BEACHHEAD

A Small Conference

44 DOC SEARLS' LINUX FOR SUITS

The Search for Terrestrial Stupidity

96 NICHOLAS PETRELEY'S /VAR/OPINION

Come Together

IN EVERY ISSUE

12 LETTERS

16 UPFRONT

48 NEW PRODUCTS

81 ADVERTISERS INDEX

INDEPTH

70 LINUX AND OPEN SOURCE IN TELECOMMUNICATIONS

What's good about being disruptive?

Ibrahim Haddad

76 SMART (SMART MONITORING AND REBOOTING TOOL)

A smarter way to monitor services.

Albert Martorell

82 A BASIC TEXT-BASED RECORDING STUDIO

You don't need a fancy GUI to create a powerful recording studio.

Matthew Geddes

86 BUILDING AND INTEGRATING A SMALL OFFICE INTRANET

Add server-side credentials to the LAMP stack.

Dave Jones



93 ADD WEB PORN FILTERING AND OTHER CONTENT FILTERING TO LINUX DESKTOPS

Station DansGuardian over incoming Web content.

Donald Emmack



48 NEW PRODUCTS

Next Month

EDITORS' CHOICE

It's Editor's choice month next month. Find out everything from our favorite distribution to our choice of word processor. There's a lot of good competition out there, and although a few of our picks were no-brainers, many of the categories deserve a long list of honorable mentions.

That's not all. We'll tell you how to use a UNIX staple called Lyx to create stellar on-line books for publishing via Lulu, and how to use two-factor authentication with ssh-agent and a USB pen-drive. Plus, Jon "maddog" Hall shares a story about the satisfaction of evangelizing FOSS.

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

Don't browse the Web

Create it

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

Try Opera.

Download the
FREE Opera browser
www.opera.com

Opera 9 - your web, your choice



HOT WIRED FOR SERVICE



A Better Way to Manage your Servers

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

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



JTLnet
web hosting specialist

Real People Real Support™

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

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

LINUX JOURNAL

Editor in Chief

Nick Petreley, ljeditor@ssc.com

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

Contributing Editors

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

Proofreader Geri Gale

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

Marketing Manager Rebecca Cassity
rebecca@ssc.com

International Market Analyst James Gray
jgray@ssc.com

Sales Coordinator Lana Newlander
ads@ssc.com

Regional Advertising Sales

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

Advertising Inquiries ads@ssc.com

Publisher Phil Hughes
phil@ssc.com

Accountant Candy Beauchamp
acct@ssc.com

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

PO Box 980985, Houston, TX 77098 USA

Editorial Advisory Board

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

Subscriptions

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

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

LINUX is a registered trademark of Linus Torvalds.



Iraq

University of Baghdad



Kosovo

Pristina



Benin

Cotonou



Denmark

Djursland



Mongolia

Ulaanbaatar



Mozambique

Maputo



South Africa

Western Cape



Iraq

Sulaimaniyah



Italy

Apennine mountains



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

Prague

MikroTik around the world
no matter what the conditions, mikrotik is routing your world. www.mikrotik.com

CIARA

DO MORE WITH LESS

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

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



Visit us

For all your computer needs

Give us a call

CIARA-TECH.COM

1-866-789-7225

IT TECH

PERSONAL CLUSTER

8 PROCESSORS 16 CORES

\$19,995

Upgrade Your System

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

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

- 150 Gigaflops computing power.
- Up to 128GB¹ FBD DDR2 533/667MHz.
- Built in 8 or 16 Port Gigabit Switch.
(Optional 8 Port InfiniBand Switch)
- Up to 8 Serial ATAII HDD.
- Integrated KVM/USB Switch.
- Plugs into one single 110/220V outlet.
- Convenient desk-side form factor.



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



Monarch Computer.com™



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



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



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

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

FREE INSTALLATION SETUP & TESTING BY CERTIFIED TECHS

FREE TECH SUPPORT!

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

AMD Motherboard Combos

FREE SETUP and TESTING

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

Only \$210

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

Only \$445



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

Only \$1,440

Supermicro H8DAE-2 PCI-E/SATA/Reg-DDR2 w/2 x AMD Opteron™ 2212 (Dual-Core) & Heatsinks (Socket F)

Only \$1,165



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

Only \$2088

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

AMD Barebone Systems

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



Antec Sonata II Piano Black Quiet ATX 2.0 Mid-Tower w/ 450w PS ASRock AM2NF6G-VSTA AMD Athlon™ 64 X2 3800+ (Socket AM2)

Starting @ \$344



Monarch Hornet Pro (Silver) w/270W PS Asus M2NPV-VM AMD Sempron™ processor 3000+ (Socket AM2)

Starting @ \$344

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

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

Starting @ \$1,325

SuperMicro (Barebone) AS1040C-8-B 1U 4-Way Dual-Core Opteron W/Hotswap SCSI & 1000W PS (Black) Supports up to 2 x AMD Opteron™ Dual-core 800 series (940) Processors

Starting @ \$1,799



Lowest Prices on AMD!

www.monarchcomputer.com

Components and Upgrades

1000s of In-Stock PC Parts and Accessories

AMD Athlon™ 64 X2/FX Dual Core Retail CPUs (939 and Socket AM2)
AMD Athlon™ 64 X2 3800+ (\$12K per core) \$148.99
AMD Athlon™ 64 X2 4200+ (\$12K per core) \$188.99
AMD Athlon™ 64 X2 4600+ (\$12K per core) \$265.00
AMD Athlon™ 64 FX-62 (1MB per core) \$829.99

AMD Athlon™ 64/FX Single Core Retail Box CPUs (939 and Socket AM2)
AMD Athlon™ 64 3500+ (\$12K per core) \$87.99
AMD Athlon™ 64 3800+ (\$12K per core) \$111.00

Dual-Core AMD Opteron™ Processors OEM CPUs (939 pin)
AMD Opteron™ 165 1.8GHz \$219.99
AMD Opteron™ 170 2.0GHz \$353.00
AMD Opteron™ 175 2.2GHz \$455.99
AMD Opteron™ 180 2.4GHz \$558.99
AMD Opteron™ 185 2.6GHz \$745.99

AMD Opteron™ OEM CPUs (939)
AMD Opteron™ 146 2.0GHz (939) \$114.00
AMD Opteron™ 148 2.2GHz (939) \$200.00
AMD Opteron™ 154 2.8GHz (939) \$560.00

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

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

AMD Opteron™ OEM CPUs (940)
AMD Opteron™ 246 2.0GHz \$158.00
AMD Opteron™ 248 2.2GHz \$203.00
AMD Opteron™ 250 2.4GHz \$307.00
AMD Opteron™ 252 2.6GHz \$441.00
AMD Opteron™ 254 2.8GHz \$668.00
AMD Opteron™ 256 2.8GHz \$825.00
AMD Opteron™ 258 2.8GHz \$977.00
AMD Opteron™ 260 2.8GHz \$1,130.00
AMD Opteron™ 262 3.0GHz \$1,469.00

Dual-Core AMD Opteron™ Processors (OEM)
AMD Opteron™ 265 1.8GHz \$270.99
AMD Opteron™ 270 2.0GHz \$395.00
AMD Opteron™ 275 2.2GHz \$540.00
AMD Opteron™ 280 2.4GHz \$720.00
AMD Opteron™ 285 2.6GHz \$918.00
AMD Opteron™ 286 1.8GHz \$677.00
AMD Opteron™ 270 2.0GHz \$847.00
AMD Opteron™ 275 2.2GHz \$1,130.00
AMD Opteron™ 280 2.4GHz \$1,469.00
AMD Opteron™ 285 2.6GHz \$2,085.00



150439 Western Digital Caviar SE 320 GB SATA 8MB Cache 7200RPM (WD3200JD) \$129.55



150353 Western Digital Raptor X 150 GB SATA 16MB Cache 10K RPM (WD1500AHFD) \$279.99



100338 Enermax Liberty ELT500AWT ATX 2.2 500W SLI Power Supply \$99.69



101208 OCZ GameXStream 700W ATX 2.2/EPS/SLI Power Supply \$149.99



1501141 3Ware Escalade 9500S-8M1 8 Port SATA RAID Controller w/Battery Backup \$443.85



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



141506 1 GB (2 pcs 512) DDR2 (667) PC2-5300 Patriot (PDC21G5300LLK) \$199.99



141801 512 MB DDR (400) PC-3200 Buffalo Select (DD4002-512/BR) \$69.99



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

NEW! Section 508 Documentation Available for our PCs!

Educational and Government POs Welcome.

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



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



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

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



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

Improving Direct Connect Architecture

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

Extending the Lead in x86 Virtualization

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

Reducing Total Cost of Ownership (TCO)

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



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

"BOTTOM LINE: MUST BUY"

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



Jason Perlow
 Linux Magazine
 April 2005

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

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



◀◀ **LOOK FOR THIS SYMBOL
 ON LINUX COMPATIBLE
 MONARCH SYSTEMS!**

The Official Linux Journal Web Server



GET QUICK QUOTES

Online or by phone:

www.monarchcomputer.com

1-800-611-0875

Paypal - Visa - Mastercard - Discover - AMEX

Monarch Empro[™] 2

4-Way 3U UNIWIDE 3546VA

Custom Server

Ask for Part #: 80488



Configuration Starting @
\$6,505!

The AMD Opteron[™] processor with Direct Connect Architecture scales from 1P/2-core up to 8P/16-core across a single industry-standard platform without external logic, allowing for maximum versatility and lowering overall system cost.

SELECTED COMPONENTS:

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

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

Supports up to 4 (8000 Series)
 Next-Generation AMD Opteron[™]
 Processors

Up to 128 GB of DDR2 Memory

1 Configuration review

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

2 Production scheduling

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

3 Board test

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

4 Build stage

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

5 Burn in diagnostic

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

6 Software load

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

7 Quality control

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

8 Packing & shipping

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

Monarch Empro[®] 2

**Custom 2U
 Rackmount SATA
 Server**

Ask for Part #: 80669



AVAILABLE COMPONENTS:

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

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

letters



Parallel Computing

Your column in the September 2006 issue [see Nicholas Petreley's "Parallel Is Coming into Its Own"], and the issue itself, was inspiring, and I found myself blogging away on the topic: (bitratchet.prweblogs.com).

—

Jed

When Good Enough Is Good Enough

Great comments by Dave Taylor [see Dave's Work the Shell column in the September 2006 issue]. I think his Blackjack script exercise was perfect for the large audience he addressed, no matter what some purists think about the endless pursuit of "perfection" in an imperfect world.

Those of you as ancient as I may recall something Gerald Weinberg passed along in *The Psychology of Computer Programming* (ISBN 0-442-29264-3): "...it is often and truly said that 'any program that works is better than any program that doesn't'" (p. 17 under section "Specifications" in my version).

Give 'em hell, Dave. You're right on the money in my book.

—

Harold Stevens

64-Bit JMP for Linux

FYI: in the September 2006 issue, Erin Vang of SAS states that R is the only statistical software available on the Linux desktop other than SAS's JMP product. However, as I'm surrounded by folks who use SAS's main competition, I recently went looking for open-source tools that might work well with SPSS (Statistical Package for the Social Sciences). My search turned up the GNU PSPP page (www.gnu.org/software/pspp), which claims:

PSPP is a program for statistical analysis of sampled data. It interprets commands in the SPSS language and produces tabular output in ASCII, PostScript or HTML format.

PSPP development is ongoing. It already supports a large subset of SPSS's transformation language. Its statistical procedure support is currently limited, but growing.

Although perhaps not as far along as other efforts, there is at least one package other than R. (I cannot comment much on how well PSPP compares to JMP or R, as I have only recently installed it and have never worked with either R or SAS software.)

—

Kevin Cole

The Dark Age of Linux Journal

First, let me make it quite clear that I have no intention of canceling my subscription. Being a reader, collector and subscriber of *LJ* since its very beginning is a honour that I will not give up that easily. But after reading an extensive letter of support in last month's issue, I felt compelled and obligated to write an as long (or longer) letter pointing to the fact that this is the worst age ever of this magazine. An age, for example, when Marcel Gagné's articles are not any longer the by-far-worst articles in any issue.

In that regard, up to a few months ago—before these Dark Ages of *LJ*—reading "Monsieur" Gagné's articles was a simple exercise of skipping the first two annoying, dull and repetitive paragraphs of his every article. From what it seemed, Mr Gagné uncreatively cut and pasted ad nausea his same little wine cellar story from previous articles. Fair enough—all one had to do was skip straight to the third paragraph, where the "good" (or at least the "better") part could be found.

Nowadays, to find a "better part" of his article, one must skip two or three *pages* and usually a comfort happens only if one can find a good advertisement—that is, not necessarily will he/she find comfort in the next also-bad article.

But Monsieur Gagné's articles have never been the chief car of the magazine. Such role is more reasonably expected from, for example, Jon "maddog" Hall, whose article this month could only be more patronizing than it is offensive to a certain "unknown" Portuguese-speaking country. Apparently, Mr Hall has visited many countries in the world, but he hasn't learned much about them, and he still belittles their inhabitants as uncivilized, uncultured and almost retarded people. In fact, I dare not ask which country he is referring to in his article out of fear that it may be the one where I was born. Nevertheless, the article in question was so childish and the dialogue reproduced therein was so painfully disconnected, pointless and senseless that I may now finally understand the reason for Jon Hall's middle (nick) name.

When things seemed bad enough, I found Dave Taylor's excuses on why his codes are so badly inefficient and yet that one should still buy or read his books and articles. In a pathetic attempt to justify himself and his apparently highly criticised lack of programming skills, Mr Taylor went over and over arguing that being a bad programmer and trying to find the easy way out is "okay"—as long as you make the proper citations, as he did in his cheating episode at UCSD.

His attempt to justify the unjustifiable could only be as degrading to oneself as the Chief Editor's, Nick Petreley, constant rebuttals to the now-so-common letters of criticisms. After all, a Chief Editor who spends his time and talent (!) to write notes in defense to what he had already defended in the

Enterprise Open Source Solutions for Demanding Data Center Environments

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

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

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

Configurable Reliable Affordable



AMD
Solution Provider
PLATINUM



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

first place (in the original article) only shows a pattern of patent and spread unpreparedness of the current staff at LJ.

I cannot really expect you to publish this letter, and I can only hope that you won't publish its parts in a distorted way in which I may sound dull and unprepared. However, I would be happy to know that my words above made you think, at least for a brief, unexpected moment. That my criticism made you (and the others to whom I am Cc'ing this message) re-evaluate what can be wrong in the magazine's new direction.

I, as anyone else, cannot assign all the blame of the current errant trend to one single person. However, when people waste pages of the magazine defending themselves—as Mr Taylor did and the CE frequently does—one starts to wonder about the coincidence of dates between this new Dark Age and the change of personnel. Either way, I still long for the days when Mr Gagné's article would—despite the boring beginnings—concentrate on the importance for our health of breaks after long uses of the computer and the availability of many software to help with that. Instead, we now find endless reports of one silly and specific Disney-like software for that purpose. Or still, two articles in the same year talking about “cool applets for KDE”. I miss the days when Jon “maddog” Hall's stories in the magazine would justify his middle name solely because of his daring, bold and yet brilliant views of a different future for the software industry, rather than his current picturesque experiences with last-century native people of “Neverland” (at least, that is how Mr Hall seems to imagine them).

Bottom line is: I hope this magazine finds its way back to being a technically rich magazine, on which people, like me, relied to read good articles: nothing more, nothing less, nothing possibly better.

--

Guilherme DeSouza

Dave Taylor replies:

Thanks for your note and your passionate enthusiasm for the publication, Guilherme. I can appreciate your desire for a more technical publication and your perspective on our editorial content, though I don't agree with it. Linux, and, by extension, software development itself, is about far more than just the lines of code. As demonstrated by the increasingly political Open Source movement, software now is the cog in the machine of commerce and as the journal of record for the Linux community, I'm proud to help offer perspectives on both the detailed geek stuff of coding and the rest of the picture.

Jon “maddog” Hall replies:

I am a little shocked that you felt my article was “patronizing and offensive”. The scene, by the

way, is Brazil. I mention real towns in it, real places and even real people (although I sometimes substitute people from Mexico and other countries). I follow this habit from one of my favorite cartoonists of all time, Walt Kelly (Pogo), who often put the names of people he had not seen for a while in his comic strip, just to let them know he was thinking of them.

The column is supposed to impart a transferral of knowledge. Most of the time the knowledge comes from me, but I also try to bring in some of the issues from the other people. A lot of the people I am “talking to” in the magazine are younger people, whose life skills are not as vast as an older person, and this would be true in any culture. If this appears to you to be condescending to the culture, I assure you that it is not meant to be that way.

I have also had people thank me for trying to bring back to the technical and commercial world the fact that Free Software is supposed to be fun.

Finally, I chose the place and the setting because I like going there, and I like the people. I will be going to an event called OpenBeach in Florianópolis, Brazil (the setting of the Beachhead) for the fourth time this year.

Marcel Gagné replies:

I write for a very different audience than Mr DeSouza would have me address. I believe that Linux and open source is good for people, all people, including the ones who want to do cool things with their desktops. I've written six books, several hundred articles and I'm coming up on seven years of Cooking with Linux. I keep writing Cooking with Linux, complete with François and my wine cellar, because people enjoy reading it. If they didn't, I would take a different tack. With a very few exceptions (such as Mr DeSouza), I get nothing but praise for my articles.

I want everybody using Linux, not just hard-core techies. Computers aren't magic and neither is software. Sometimes I feel that if we can't reach out to the average person, explain things in simple terms whenever possible, and make it fun for them, we aren't doing our jobs right. If offering up a wine suggestion with every column makes my discussion of desktop backup solutions, multimedia jukeboxes, panel applets, desktop search engines and so forth more fun, then so be it.

Mr DeSouza has every right to express his feelings, whether I agree with them or not (and I don't), but I'm not writing for him. Apparently, none of us are.

Nicholas Petreley replies:

I'll take your advice and decline to defend the fact that I've written rebuttals in response to some critical letters.

LINUX JOURNAL

At Your Service

MAGAZINE

PRINT SUBSCRIPTIONS: Renewing your subscription, changing your address, paying your invoice, viewing your account details or other subscription inquiries can instantly be done on-line, www.linuxjournal.com/subs. Alternatively, within the U.S. and Canada, you may call us toll-free 1-888-66-LINUX (54689), or internationally +1-713-589-2677. E-mail us at subs@linuxjournal.com or reach us via postal mail, Linux Journal, PO Box 980985, Houston, TX 77253-3587 USA. Please remember to include your complete name and address when contacting us.

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

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

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

ON-LINE

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

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

Linux laptops. Supported.

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

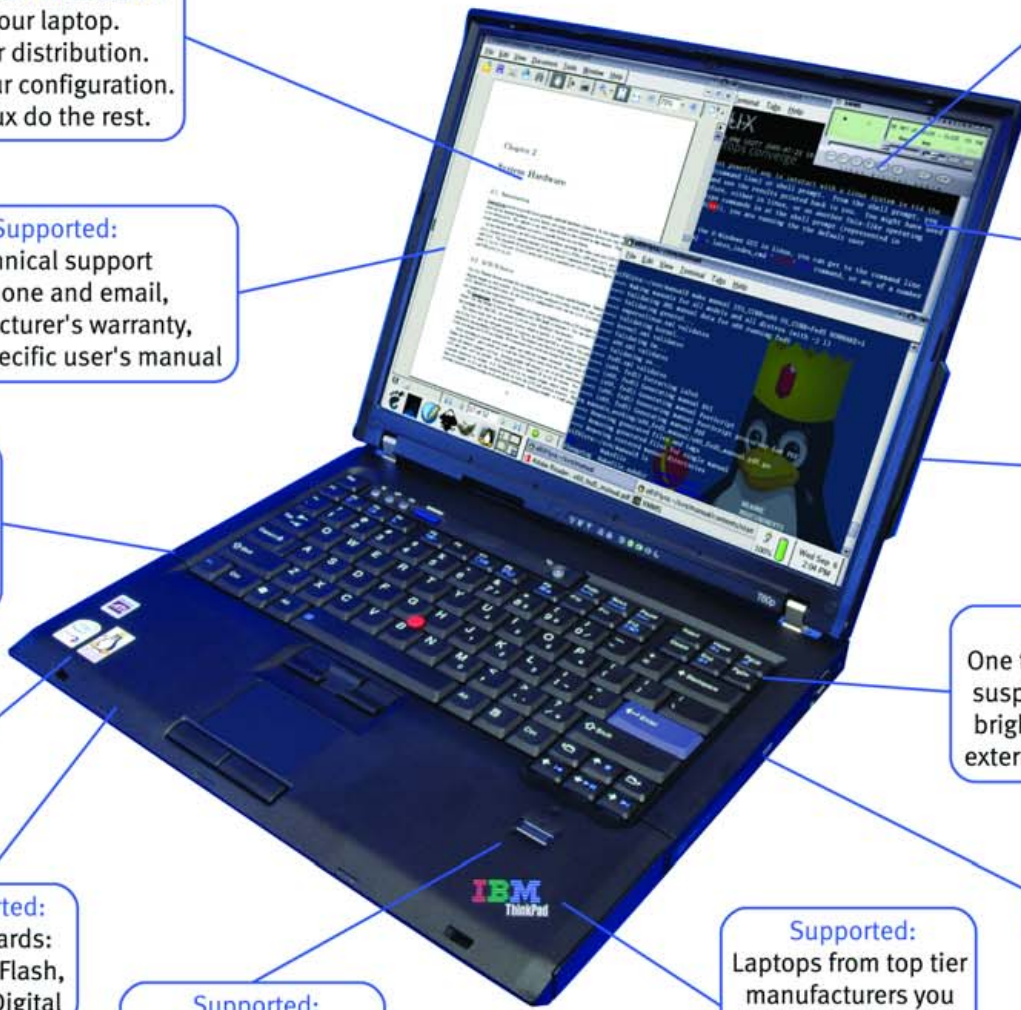
Supported:
Technical support
by phone and email,
manufacturer's warranty,
system-specific user's manual

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

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

Supported:
Media cards:
Compact Flash,
Secure Digital

Supported:
Biometric fingerprint
(GDM login with PAM)



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

Supported:
Power management:
suspend, hibernate,
processor control

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

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

Supported:
Internal optical drive:
CDRW, DVD±RW

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

Since 1999, EmperorLinux has provided pre-installed Linux laptop solutions to universities, corporations, and individual Linux enthusiasts. We specialize in the installation of the Linux operating system on a wide range of the finest laptops and notebooks made by IBM, Lenovo, Dell, Sharp, Sony, and Panasonic. We offer a range of the latest Linux distributions, as well as Windows dual boot options. We customize each Linux distribution to the particular machine it will run upon and provide support for: ethernet, wireless, EVDO mobile broadband, PCMCIA, USB, FireWire, X-server, CD/DVD/CDRW, sound, power management, and more. All our systems come with one year of Linux technical support by both phone and email, and full manufacturers' warranties apply.
Visit www.EmperorLinux.com or call 1-888-651-6686 for details.

EmperorLinux

...where Linux & laptops converge

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

Model specifications and availability may vary.



diff -u

WHAT'S NEW IN KERNEL DEVELOPMENT

The **Areca RAID driver** is likely to go into the main kernel sources soon. **Andrew Morton** and others have been working with the maintainer, **Erich Chen**, to fix some remaining issues, and this work seems to be bouncing right along. There was a small disconnect when the official merge was first proposed, as some folks hadn't realized Erich had been actively maintaining the driver or working on addressing the problems.

Jesse Huang from **IC Plus** released an **IP100A 10/100 fast network adapter driver** for inclusion, and a little corporate competition came into play. Jesse's code was a fork of the **Sundance driver code**, with only minor changes. In such cases, the going wisdom, as **Arjan van de Ven** put it, is to update the existing driver to support the additional hardware, instead of starting a whole new driver. In response to this, Jesse explained that IC Plus was keen to have the ip100a.c filename in the kernel and possibly remove sundance.c at some point. But, after conferring with his company, they decided to follow best practices and just feed their changes into sundance.c.

Al Boldi submitted a patch to make **RootFS** swappable by using tmpfs rather than ramfs for its back-end file storage. This change would allow systems with a large initrd or initramfs image not to tie up the RAM associated with that image unless it is actually in use. The idea behind Al's patch did make sense to folks as being useful for embedded systems. But, **H. Peter Anvin** pointed out that one current goal is to move initramfs initialization earlier in the boot process to include loading firmware for built-in device drivers. He wasn't sure how Al's patch would affect this plan, if at all, but he said the migration to earlier initialization had to take precedence.

Nigel Cunningham has submitted his **Suspend2** code for inclusion in Andrew Morton's -mm tree. In spite of forking the software suspend

code from **Pavel Machek** years ago, this is the first time he's actually submitted it for inclusion anywhere. A lot of users find Suspend2 to be much better than Pavel's uswsusp code, and they routinely download and apply Nigel's patch even though uswsusp is already in the kernel. Pavel, as the official software suspend maintainer, has nearly the final word on whether anyone else's suspend code gets into the kernel, and his antagonistic relationship with Nigel makes it unlikely the he would let the code through without a fight. But, his arguments against including Nigel's code have begun to ring hollow. He says the in-kernel code works just as well, but then hordes of users proclaim that no, Suspend2 works better for them. He says uswsusp is a better idea and users should just wait for it to be ready, but Suspend2 works now and has worked well for a long time. It does seem as though Nigel's code has proven itself, and without serious technical objections, it should be allowed into the kernel.

Hans Reiser is at it again, claiming that kernel developers have been standing in the way of including **Reiser4** for political reasons. Although some kernel brawls do seem to be politically motivated, Hans just doesn't have the high ground. He's repeatedly hurled attacks and insults at the kernel developers reviewing his code—to the point where several key developers now refuse to offer any more technical feedback on Reiser4. Without these reviewers, it becomes very difficult for the Reiser developers even to identify the remaining technical issues that must be addressed before the code could be included. Because Hans doesn't seem able to see how antagonistic he's been, perhaps his friends should urge him to stay out of kernel debates and let the other Reiser4 engineers speak for him. It seems to me that the same people who currently refuse to work with Hans would be happy to rejoin the effort if they didn't have to fear his attacks.

—Zack Brown

FIRST LOOK:

Sony's New mylo Handheld

mylo (for "my life online") is Sony's new competitor against the Nokia 770 in the Linux-based handheld computer category. It's a bit smaller (1 x 4.8 x 2.5 inches), has a 2.4-inch QVGA (320 x 240) LCD screen, and a retractable keyboard. Where the 770 is a rectangular tablet (with a much larger screen), the mylo has rounded corners and looks more like a mobile phone.

Like the 770, however, the mylo is not a phone, but rather supports IP telephony systems, such as Skype, which is also listed by Sony as one of its four mylo "partners". As of August 2006, the others are JiWire (for finding 802.11b Wi-Fi hotspots), Yahoo and Google (both for instant messaging and e-mail).

Perhaps most significant, from a historical perspective, is that Sony is supporting audio formats other than its own. The mylo comes with support for MP3 audio, as well as Sony's own ATRAC3 and Microsoft's WMA. It also has a built-in MPEG-4 video player. Until now, Sony has avoided making MP3 players, a category now dominated by Apple's iPod.

Files can be transferred to and from the mylo either by USB2 connections or Sony's proprietary (but common) Memory Stick removable Flash media.

Sony hasn't released any hardware specs (such as processor or speed), but among the specs it shares are 1GB internal Flash RAM, a rechargeable 3.7-volt battery (and external 6 V DC power adapter), video playing time of up to 8 hours and talk time of up to 3.5 hours.

In general, the mylo is designed to work immediately as a consumer electronics device. But, it's still a Linux-based computer. And, like the 770, it is open to application development through the Qtopia platform from Trolltech.

We will be taking a closer look at the mylo in the next few months. In the meantime, feel free to share your own experiences with the device. Write to ljeditor@ssc.com.

See www.learningcenter.sony.us/assets/itpd/mylo/prod/index.html, en.wikipedia.org/wiki/Sony and linuxdevices.com/news/NS8202297251.html for more information.

—Doc Searls

They Said It

Used to be I couldn't spell genus and now I are one.

There are a lot of computer languages out there that are doing drugs.

If there's one problem Perl is trying to solve it's that all programming languages suck.

It takes ten years to become good at being a kid. Then another ten years to become good at not being a kid.

An adult is someone who knows when to care.
—All from a speech by Larry Wall at OSCON 2006

I'm not much interested in interoperability. I want substitutability. I want to be able to throw your software out.

—Simon Phipps, talk at OSCON 2006

Universities love to include pictures of their CIOs. I have no idea why.

—Steven O'Grady, talk at OSCON 2006

There is nothing as strong and as indestructible as a mesh network. And that's what the Internet is.

—Tom Esvlin, at a Berkman meeting

What could you do with **216** cores



THINKMATE 8-WAY WORKSTATION FEATURING AMD OPTERON™ PROCESSORS

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

Better Efficiency, Greater Productivity, and Enhanced Scalability

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

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

Video
PCI Express x16 or Integrated ATI Rage XL 8MB

Storage
Up to 8 Removable Serial ATA or SCSI Hard Drives

Power Supply
1350W 3+1 Redundant Power Supplies

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

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

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

**16
Cores**

\$62,999

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

**16
Cores**

\$17,999

THINKMATE

(800) 371-1212
WWW.THINKMATE.COM

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

LJ Index, November 2006

1. Percentage of new cars sold in the US that will have a jack that works with Apple iPods: **70**
2. Number of devices other than Apple's that the new jack will work with: **0**
3. "Weekly audience" rank of NPR among radio formats*: **4**
4. "Listened to Most Often" rank of NPR among radio formats: **2**
5. "Conversion to Most Often Listener" rank of NPR among radio formats: **1**
6. Millions of people who will be watching TV on cellular handsets by 2011: **446**
7. Percentage year-on-year growth rate of mobile phone TV through 2010: **50**
8. Percentage of Koreans hooked to a broadband network: **90**
9. Projected trillions of US dollars generated by business made possible by U-Japan ("ubiquitous networked society") by 2010: **1**
10. Rank of South Korea in broadband penetration: **1**
11. Rank of Canada in broadband penetration: **8**
12. Rank of Luxembourg in broadband penetration: **19**
13. Rank of US in broadband penetration: **20**
14. Millions of Linux-based Motorola smartphones shipped in China during Q2 2006: **1**
15. Rank of Motorola among providers of cell phones in China: **2**
16. Percentage of top five mobile device vendors with a "Linux strategy": **80**
17. Number of members in OSDL's Mobile Linux Initiative: **15**
18. Year by which Linux will surpass Symbian as the top mobile OS: **2010**
19. Linux mobile OS market-share percentage at the end of 2005: **23**
20. Microsoft mobile OS market-share percentage at the end of 2005: **17**

Sources: 1, 2: Marketplace Radio | 3–5: Center for Media Research, reporting on The Media Audit | 6, 7: IMS Research, reported in *LinuxDevices* | 8, 9: *The Age* | 10–13: Point Topic, via *WebSiteOptimization.com* | 14–16: *LinuxDevices* | 17, 18: *LinuxDevices*, citing OSDL and The Diffusion Group | 19, 20: *Total Telecom*, citing The Diffusion Group

* NPR is a network and not a format, but the study treated it as a format.

—Doc Searls

Lenovo Makes the Commitment

Around two years ago, when Novell announced a corporate commitment to move the company completely to Linux-based hardware, its laptop of choice was IBM's ThinkPad. At LinuxWorld Expo and other Linux-related conferences, ThinkPads were about the only brand of laptop populating Novell booths.

Then, after IBM sold its PC division to China-based Lenovo, many wondered if the ThinkPad would survive, or if the company would pay close attention to potential Linux customers. But, ThinkPads continued to sell, now with Lenovo instead of IBM printed on their cases. Some users even began waxing enthusiastic about them. In June 2006, Cory Doctorow, the prolific writer of science-fiction books and the top-ranked BoingBoing blog, announced that he was switching after many years from Mac to Linux:

I thought about buying a MacBook Pro anyway, since they're nice computers, and they run Ubuntu, but after pricing them out, I realized that I could get a lot more bang for my buck with a Lenovo ThinkPad T60p. If I'm not going to run the Mac OS, why spend extra money for Apple hardware? I ordered the machine last weekend, loading it to the max with two 120GB hard drives, 2GB of RAM, and the fastest video card and best screen Lenovo sells. It was still cheaper than a Mac, even though Lenovo makes me pay for a copy of Windows XP that I plan on pitching out along with the styro-foam cutouts and other worthless packaging.

With that kind of writing on the wall, something big was bound to happen. And, at the latest LinuxWorld Expo [August 2006], it did. Lenovo revealed that it would make the first Linux-based ThinkPad "mobile workstation". It will come with Novell SUSE Linux Enterprise 10, on a ThinkPad T60p, which is built around Intel's new 2.3GHz Core Duo T2700 chipset. According to Novell PR honcho Bruce Lowry, the new offering is the product of a joint effort between Lenovo, Novell and Intel engineers.

For the last three years, most of my Linux life has been on a ThinkPad T40, most recently running Novell's SUSE Linux desktop. It's been good, but it's also been a hermit crab. You can tell by the "Access IBM" button that works only if the machine is running Windows. Well, on the Linux-equipped ThinkPad T60p, that button gets you the Lenovo Help Center, which covers ThinkVantage Technologies, drivers, basic Linux configuration and hardware issues. Novell handles core operating system issues.

Both companies are working to make sure media runs well on the machine too. By the end of this year, an upgraded RealPlayer will play Windows media files, the company says. Lenovo also says the current Helix Banshee player in SUSE Linux Enterprise Desktop (SLED) 10 is the only Linux software that allows encoding of MP3 audio files and burning audio CDs.

I'm looking forward to trying out this new configuration. (As a notoriously clumsy user, I expect to give the help desks a workout.) Meanwhile, look for Cory Doctorow's Mac-to-Ubuntu migration account in an upcoming issue of *Linux Journal*.

—Doc Searls

USER FRIENDLY by J.D. "Illind" Frazer



LINUX JOURNAL EDITION

AMD® Dual Core to 16 Processors!



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

Powerful & Quiet w/DDR2 800!

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



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

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

Xi MTower 64 x2 4200+ \$1,499
Xi MTower 64 x2 FX-62 \$2,249

Xi MTower 2P64 2x2212 \$2,749
Xi MTower 2P64 2x2220 \$5,299

Xi NetRAIDER 64 2x2212 \$ 5,499
Xi NetRAIDER 64 4x8220 \$16,999



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

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

@XiComputer.com

Intl:(949)498-0858

Fax:(949)498-0257

www.xicomputer.com

1-800-432-0486



REUVEN M. LERNER

Beginning Ajax

How to put the A (asynchronous) in Ajax.

Using Web standards, such as the DOM (document object model) and CSS (cascading stylesheets), Ajax applications can approach the usability, friendliness and instant feedback that people expect from desktop applications.

Many programmers, myself included, have long seen JavaScript as a way to change the appearance of a page of HTML dynamically or to perform relatively minor tasks, such as checking the validity of a form. In the past year, however, JavaScript has emerged as a major force for application developers, providing the infrastructure for so-called Ajax applications.

Before JavaScript, there was a one-to-one correspondence between user actions and the display of HTML pages. If a user clicked on a link, the currently displayed page disappeared and was replaced with another page of HTML. If a user submitted an HTML form, the contents of that form were submitted to a program on the Web server, and the content of the server's response was then displayed in the browser, replacing its predecessor. In traditional Web applications, server-side programs handle the bulk of user input and also build any dynamically generated Web pages the user might see.

Ajax applications redistribute the load, putting a greater emphasis on client-side JavaScript. In an Ajax application, many server-side programs do indeed produce complete pages of HTML, which are then displayed in their entirety in a Web browser. But many other server-side programs produce small snippets of XML-formatted data. This data is both requested and used by client-side JavaScript to modify and update the current HTML page without having to refresh or replace it. Using Web standards, such as the DOM (document object model) and CSS (cascading stylesheets), Ajax applications can approach the usability, friendliness and instant feedback that people expect from desktop applications.

This month, we continue exploring client-side JavaScript and Ajax, which we began during the past few months. Last month's column looked at a user-registration application for a Web site. Although the actual registration took place in a server-side program, we looked at ways in which we could provide an Ajax-style warning for registering users who wanted a user name that was already taken. Sure, we could have the server-side registration program check to see whether the user name had been taken already, but that would require refreshing the page, which also requires a delay.

The solution we implemented last month was fine from the user's perspective (especially if the user has somewhat Spartan tastes in design), but it solved the problem in a very non-Ajax way—by hard-coding the user names in a JavaScript array and then looking for the desired new user name in that array. This approach has a number of large problems associated with it, starting with the fact that the full list of user names is available to anyone looking at the HTML source and ending with the fact that the array will become unwieldy and cumbersome over time, taking an increasingly long time to download and search through as the number of registered users grows.

We can avoid these problems by using an Ajax-style solution. Rather than hard-code the list of user names in the JavaScript, and instead of having the server-side program produce a full list of user names, perhaps we could simply send a request to the server, checking to see if the requested user name is already taken. This will result in relatively fast download and reaction times, in a

cleaner application design and in an extensible application.

This month, we take the Ajax plunge, modifying the server- and client-side programs we wrote last month to retrieve user names via an asynchronous request from the server. In producing this application, we will see how relatively straightforward it can be to create an Ajax application or to integrate Ajax functionality into a traditional Web application. By the end of this article, you should understand how to create the client and server sides of an Ajax application.

Making an Ajax Call

The technology that makes much of Ajax possible is JavaScript's XMLHttpRequest object. Using this object, a JavaScript function can make HTTP requests to a server and act on the results. (For security reasons, HTTP requests made by XMLHttpRequest must be sent to the server from which the current Web page was loaded.) The HTTP request may use either the GET or POST method, the latter of which allows us to send arbitrarily long, complex content to the server.

Most interesting, and at the core of many Ajax paradigms, is the fact that XMLHttpRequest may make its HTTP requests synchronously (forcing the browser to wait until the response has been completely received) or asynchronously (allowing the user to continue to use the browser window as it downloads additional information). Ajax applications typically use asynchronous calls. This allows different parts of the Web page to be updated and modified independently of one other, potentially responding simultaneously to multiple user inputs.

Ideally, we would be able to create an instance of XMLHttpRequest with the following JavaScript code:

```
var xhr = new XMLHttpRequest();
```

Unfortunately, life isn't that simple. This is because many people use Internet Explorer as their primary browser. IE does not have a native XMLHttpRequest object, and thus it cannot be instantiated in this way. Rather, it must be instantiated as:


```
var xhr = new ActiveXObject("Msxml2.XMLHTTP");
```

But wait! There are also some IE versions that require a slightly different syntax:

```
var xhr = new ActiveXObject("Microsoft.XMLHTTP");
```

How are we going to handle these three different ways of instantiating XMLHttpRequest? One way is to use server-side browser detection. It is also possible to use client-side browser detection. But the most elegant method I have seen to date comes from *Ajax Design Patterns*, a new book by Michael Mahemoff (published by O'Reilly Media). Mahemoff uses JavaScript's exception-handling system to try each of these in turn until it works. By wrapping our three different instantiation methods in a function, and then assigning the value of our xhr variable to whatever the function returns, we can give

The Industry Leader for Server Appliances



Custom server appliances or off the shelf reference platforms,
built with your image and software starting under \$1,000.
From design to deployment, we handle it all.

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

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

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



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

our application cross-platform compatibility:

```
function getXMLHttpRequest () {
  try { return new ActiveXObject("Msxml2.XMLHTTP"); } catch(e) {};
  try { return new ActiveXObject("Microsoft.XMLHTTP"); } catch(e) {}
  try { return new XMLHttpRequest(); } catch(e) {};
  return null;
}

var xhr = getXMLHttpRequest();
```

After executing the above code, we can be sure that `xhr` is either null (indicating that all attempts to instantiate `XMLHttpRequest` failed) or contains a valid instance of `XMLHttpRequest`. Once instantiated, `XMLHttpRequest` is compatible across browsers and platforms. The same methods thus will apply for all systems.

Listing 1. ajax-test.html

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

  <script type="text/javascript">
    function getXMLHttpRequest () {
      try { return new ActiveXObject("Msxml2.XMLHTTP"); } catch(e) {};
      try { return new ActiveXObject("Microsoft.XMLHTTP"); } catch(e) {}
      try { return new XMLHttpRequest(); } catch(e) {};
      return null;
    }

    function parseHttpResponse() {
      alert("entered parseHttpResponse");
      if (xhr.readyState == 4) {
        alert("readystate == 4");
        if (xhr.status == 200) {
          alert(xhr.responseText);
        }
        else
        {
          alert("xhr.status == " + xhr.status);
        }
      }
    }

    var xhr = getXMLHttpRequest();
    alert("xhr = " + xhr);
    xhr.open("GET", "atf.html", true);
    xhr.onreadystatechange = parseHttpResponse;
    xhr.send(null);

  </script>
</head>
<body>
  <h2>Headline</h2>
  <p>Paragraph</p>
</body>
</html>
```

The most common method to call on `xhr` is `open`, which tells the object to send an HTTP request to a particular URL on the originating server. A call to `xhr.open` looks like this:

```
xhr.open("GET", "foo.html", true);
```

The first parameter (`GET`) tells `xhr.open` that we want to use the HTTP `GET` method. The second parameter names the URL that we want to retrieve; notice that because we must connect to the originating server, the initial protocol and host-name part of the URL is missing. The third parameter indicates whether the call is asynchronous (`true`) or synchronous (`false`). Almost all Ajax applications pass `true`, as this means that the browser doesn't freeze up while it is waiting for the HTTP response. This ability to make asynchronous HTTP requests is central to the magic of Ajax. Because the HTTP request doesn't affect the user interface and is handled in the background, the Web application feels more like a desktop application.

The call to `xhr.open()` does not actually send the HTTP request. Rather, it sets up the object so that when the request is sent, it uses the specified request method and parameters. To send the request to the server, we use:

```
xhr.send(null);
```

`XMLHttpRequest` does not return the HTTP response whoever calls `xhr.send()`. This is because we are using `XMLHttpRequest` *asynchronously*, as specified with the `true` value to `xhr.open()`. We cannot predict whether we will get results in half a second, five seconds, one minute or ten hours.

Instead, we tell JavaScript to invoke a function when it receives the HTTP response. This function will be responsible for reading and parsing the response and then taking appropriate action. One simple version of the function, which I have called `parseHttpResponse`, is as follows:

```
function parseHttpResponse() {
  alert("entered parseHttpResponse");
  if (xhr.readyState == 4) {
    alert("readystate == 4");
    if (xhr.status == 200) {
      alert(xhr.responseText);
    }
    else
    {
      alert("xhr.status == " + xhr.status);
    }
  }
}
```

`parseHttpResponse` is called when the HTTP response to our Ajax request comes in. However, we have to make sure that the response contents have completely arrived, which we do by monitoring `xhr.readyState`. When that equals 4, we know that `xhr` has received the complete response. Our next step is then to check that the response had an HTTP "OK" (200) code. After all, it is always possible that we got a 404 ("file missing") error from the server, or that we failed to connect to the server at all.

To tell JavaScript we want to invoke `parseHttpResponse` when our HTTP request returns, we set the `onreadystatechange` attribute in our `XMLHttpRequest` object:

```
xhr.onreadystatechange = parseHttpResponse;
```


Expand Your Computing Density



AMAX HPC blade servers leverage leading edge dual core processing technology to provide optimized computing solutions with better thermal efficiency, higher performance per watt, higher operational effectiveness, deployment flexibility, simplified management, SAS storage, various interconnect options, and Virtualization capabilities for deployment by many data centers around the world.


Our ISO-9001 certified Manufacturing facilities and highly skilled engineers provide customers with reliable, secure, cost-effective, and flexible computing solutions.

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



Expand Your HPC Flexibility

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



AMAX Global Services Also Available!
Toll Free: 1.866.626.2948 ext.1106
WWW.AMAX.COM

If you've ever wondered how hard it is to perform an Ajax call, you now can see that it's relatively simple.

Finally, after we can be sure that we have received the response and that all is well, we can grab the text of the response with the `xhr.responseText` method. Our XMLHttpRequest can return its response either as a text string (as here) or as an XML document. In the latter case, we then can use the DOM to navigate through it, much as we would do with a Web page.

Of course, an actual Ajax application would not issue an alert at every step of its execution and would probably do something more useful—perhaps changing some text, adding or removing some nodes from the document tree or changing part of the document's stylesheet. Nevertheless, you can see this code in action in Listing 1 (`ajax-test.html`).

Note that `ajax-test.html`, although simple, is a fully working Ajax program. In order for it to work, you need to have a file named `atf.html` in the DocumentRoot directory of your Web site. (Otherwise, you will get an HTTP response code of 404.) If you've ever wondered how hard it is to perform an Ajax call, you now can see that it's relatively simple.

Adding Ajax to Registration

Now that we have seen how an Ajax program works, let's use this knowledge to modify the registration program that we built last month. Our old registration page defined a list of user names in the JavaScript. If the user's requested user name was a member of that list, we alert the user to the error and forbid the user from actually registering.

I won't describe all of the problems with this approach, as there are many. As a simple alternative, what if we were to use Ajax to retrieve the list of user names? That way, we could be sure that the list was up to date.

What if, instead of having the array contents hard-coded, we were to download them from a Web page on the server? (This is admittedly not as sophisticated as getting a yes or no answer to a specific user name; we will get to that functionality in next month's column.) If the Ajax-retrieved list of user names was generated dynamically, we could have it grab appropriate data from the database and then return an XML document that easily could be turned into an array. To make the example easier in this month's column, we don't use a dynamic page, but rather a static one. However, if you have done any server-side Web programming in the past, you probably will understand how to take our file, `usernames.txt` (Listing 2), and turn it into a dynamic page.

A registration page that follows this principle is shown in Listing 3. That file, `ajax-register.html`, is similar to the registration form we created last month. In last month's non-Ajax version, we defined an array (`usernames`). We then defined a `checkUsername` function that is invoked by the `onchange` handler for the username text field. This had the effect of invoking `checkUsername` when the user completed the user name. If the requested user name was in the `usernames` array, the user was given a warning, and the submit button was disabled. Otherwise, the user was able to submit the form to the server-side registration program, presumably as a first step to participating in the site.

To turn last month's registration page into an Ajax-style one, we modify the `checkUsername` function, which is invoked when the user finishes entering his or her requested user name. Instead of defining the `usernames` array, we instead have `checkUsername` fire off an Ajax request to the server. Unlike last month's non-Ajax version, this is all that `checkUsername` will do. The updated function looks like this:

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

As you can see, our function is requesting the file `usernames.txt` from the server. When `xhr`'s state changes, we ask to invoke the `parseUsernames` function. It is in this function that we have put the serious logic, first turning the retrieved file contents into an array:

```
var usernames = [ ];

if (xhr.readyState == 4) {
  if (xhr.status == 200) {
    usernames = xhr.responseText.split("\n");
  }
}
```

Here, we see the standard Ajax pattern repeated from the previous example: wait for `xhr.readyState` to be 4, and then check that `xhr.status` (the HTTP response status code) is 200. At that point, we know we have received the contents of `usernames.txt`, which (as you can see from Listing 2) contains the existing user names, one user name per line. We use JavaScript's `split` function to turn this into an array, which we assign to `usernames`.

From this point on, we can reuse the logic from last month's non-Ajax version, first grabbing the various node IDs from the page, using DOM methods:

```
var new_username = document.forms[0].username.value;
var found = false;
var warning = document.getElementById("warning");
var submit_button = document.getElementById("submit-button");
```

Then, we check to see if the requested user name is in our array:

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


If the user name is found in the list, we issue a warning at the top of the page. Otherwise, we clear out any warning that might be there:

```
if (found)
{
  setText(warning, "Warning: username '" + new_username + "' was taken!");
  submit_button.disabled = true;
}

else
{
  removeText(warning);
  submit_button.disabled = false;
}
```

Listing 2.
`usernames.txt`

```
abc
def
ghi
jkl
mno
pqr
stu
vwx
yzz
```



Boost Performance on Linux Clusters.

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

Intel® C++ and Fortran Compilers for Linux*

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

Intel® Cluster Toolkit for Linux

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

▪ Intel® MPI Library

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

▪ Intel® Math Kernel Library Cluster Edition

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

▪ Intel® Trace Analyzer and Collector

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

▪ Intel® MPI Benchmarks

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

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

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

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

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

Come visit us at SC06 in booth 1523.



Listing 3.

ajax-register.html

```

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

  <script type="text/javascript">
function getXMLHttpRequest () {
  try { return new XMLHttpRequest("Msxml2.XMLHTTP"); } catch(e) {};
  try { return new XMLHttpRequest("Microsoft.XMLHTTP"); } catch(e) {};
  try { return new XMLHttpRequest(); } catch(e) {};
  return null;
}

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

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

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

var xhr = getXMLHttpRequest();

function parseUsernames() {
  // Set up empty array of usernames
  var usernames = [ ];

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

}

// Get the username that the person wants
var new_username = document.forms[0].username.value;
var found = false;
var warning = document.getElementById("warning");
var submit_button = document.getElementById("submit-button");

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

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

}

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

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

```

Now, is this a good way to handle the checking of user names? Not really—although now that we have the basic Ajax logic in place, we can modify it slightly to be more efficient and secure.

One problem is that the list of user names is in a static file. Perhaps our server is running a cron job that creates usernames.txt on a regular basis, but that seems a bit silly when we can instead use a server-side program to query the database dynamically. Switching from a static file to a dynamic page thus seems like a good idea, if only for performance reasons.

There are security reasons as well. As with last month's version, we are downloading the entire list of user names to the user's browser. This means that a potentially malicious user would have access to all of the user names and would be able to poke through them, either with the intention of trying to break into the site or spam the users.

One potential downside of using Ajax for this type of check is the speed issue. As I indicated previously, the core of Ajax is its asynchronous nature, which means that we cannot know how long it will take for the server to respond to our query. In my simple tests, the round trip from my browser to my server and back was nearly instantaneous, and it provided me with useful feedback right away. On a more heavily loaded server, or with a more sophisticated database query, or if users have slow Internet connections, asynchronous calls might begin to feel sluggish. That said, even the worst Ajax function will likely be faster than a page refresh, because of the reduced overhead that is involved.

Conclusion

This month, we finally begin to use Ajax in an application. We see here how it is possible to take some existing JavaScript code and break it apart into two functions: one that invokes the Ajax call and the other that handles the parsing of data when the call receives a response.

However, we also see that there are security and efficiency problems with this approach. A better technique would be to send only the requested user name in the Ajax call and get a simple yes or no answer from the server, indicating whether the user name had been taken already. Next month, we will do just that, using an Ajax POST query instead of our GET query from this month, and replacing usernames.txt with a server-side program that works in conjunction with our Ajax call. ■

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

Further Reading

There has been an explosion of books and articles about Ajax programming in the last year, and I am slowly making my way through many of them. Two of the best that I've read are both published by O'Reilly. *Head Rush Ajax* is aimed at beginners and teaches the introductory material in a fun, effective way. *Ajax Design Patterns*, which I mentioned earlier in this article, is probably my favorite Ajax book so far (despite its design and editing, which aren't up to the usual O'Reilly standards). This latter book is a good introduction to the subject for experienced Web developers.

The Ajaxian.com Web site has a large number of links, tutorials and articles having to do with Ajax development on a variety of different platforms. If you're interested in Ajax development, it's worth keeping this site in your RSS reader or bookmarks.

**ASA
COMPUTERS**
www.asacomputers.com
1-800-REAL-PCS

Hardware Systems For The Open Source Community—Since 1989

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

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

AMD Opteron™ Value Server- S795

1 U 14.3" Deep
AMD Opteron 140 1M Cache
1 GB DDR ECC REG PC-3200
1 of 2 40GB SATA Drive
2 X 10/100/1000 NIC
Options: CD, FD, or Second Drive, Raid
ADD Your Logo



ISCSI Dual AMD Opteron™ 1U to 8U, Call for Pricing

1TB to 30TB of iSCSI Storage
Dual AMD Opteron 246
1 GB DDR ECC REG PC-3200
Dual GigE LAN
Redundant PS, Hot-Swap Drives
RAID Options, RAID 5, 10, 50
More Customization is available



1U SCSI Quad AMD Opteron™ Starting @ \$2850

1 of 4 AMD Opteron 848
2 GB DDR ECC REG PC-3200
1 of 3 36GB SCSI Drive
2 GigaE, CD, FD,
Remote Management Card (IPMI)



30TB AMD Opteron™ Storage Solution- Starting @ \$26,395

30TB SATA Storage in 8U
Includes all Raid Cards, Raid 5, 10
Dual AMD Opteron 246
2 GB DDR, ECC REG PC-3200
Dual GigE, FD, CD



Your Custom Appliance Solution

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



Custom Server, Storage, Cluster, etc. Solutions

Please Contact us for all type of Storage solutions, NAS, DAS, iSCSI, Fiber RAID, SATA, SAS
*Free shipping on selected servers and all notebooks.



2354 Calle Del Mundo, Santa Clara, CA 95054

www.asacomputers.com

Email: sales@asacomputers.com

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

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



MARCEL GAGNÉ

The Dynamic Web: for Those Who Like to Watch

Nothing says dynamic like video, and some people out there are doing some amazing things. Of course, I know some people would rather sit back and watch all this action than make it themselves. So grab your remote, sit back and enjoy your wine!

One of the coolest programs I've run across in a while makes this whole mess of trying to find great Internet videos just that much easier.

Yes, François, it does seem that way sometimes. Despite all the great hopes that the Web would become a place where the world could interact and share knowledge, that knowledge does occasionally tip toward entertainment. The epitome of dynamic Web development seems to have culminated in a new video delivery system. Of course, you can learn a great deal from videos, and many sites allow the people who visit to comment and discuss the ideas expressed in those videos. It's true that some of what we find out there is effectively cut and pasted from television, but some people are taking advantage of this relatively inexpensive video delivery medium to stretch their creative muscles and share their ideas with the world. What it does show, in my opinion, is that there is an amazing amount of talent out there. This technology not only enables those people to reach out to the world, but it also enables the watchers to experience fresh, new talent.



Figure 1. You'll love Democracy's integrated channel guide. Find the shows that interest you, and add the channel with a click.

I see that our guests approach, François. Pay attention to today's menu. I've got something that takes this whole Internet television concept to a new level. Ah, here they are. Welcome, *mes amis*, to *Chez Marcel*, home of fine Linux cuisine and exquisite wines. François and I were just discussing the explosion of video on the World Wide Web and what it means for both content creators and video consumers. Please, sit down while I send my faithful waiter to fetch the wine. I believe we still have two cases of the 2000 Château La Tour Blanche in the cellar. Please, bring it back quickly.

There's a lot of content out there, and Internet video broadcasting (or vodcasting) is creating a busy landscape, sometimes confusing to navigate. One of the coolest programs I've run across

in a while makes this whole mess of trying to find great Internet videos just that much easier. Called Democracy, this is an open-source Internet television watching program with an interesting mandate at its core. The group behind Democracy is a not-for-profit organization that calls itself the Participatory Culture Foundation. This group, like many others I'm sure, is concerned about the fact that so much of our media is controlled and filtered through a handful of large corporations. It feels that the best hope for dealing with this kind of centralized editorial and media control is to support an open standards, participatory Internet TV format. Part of its solution is the Democracy Internet TV media player.

Here's how it works. The program lets you find, download, record, manage and watch Internet television programs. Those of you who have a TiVo at home will understand the beauty of this concept. By default, downloaded videos are saved for five days, at which point they are automatically deleted. If you want to save them permanently, you do have that option. Democracy also plays in full screen, so you can take advantage of that 21-inch flat-panel screen in front of you. With an integrated channel guide, a community-based rating system and associated publishing tools (so you can get in on the action), there's a lot more to Democracy than just watching TV. Because you'll more than likely start out watching content, I'll tell you all about that in a moment. First, though, I think François has returned. Please, *mon ami*, make sure everyone's glass has been filled.

To get your copy of Democracy, pay a visit to its Web site (see the on-line Resources). The site offers packages for Ubuntu, Fedora and plain-old Debian. Source is also available if you can't find a package for your particular distribution. Installing Democracy TV is no big deal, but I should warn you that it does require Mozilla and the associated packages. I make a point of mentioning this because many of us are now running Firefox instead of Mozilla. If you are downloading Mozilla as a package (or packages), make sure you have the development and PSM packages as well.

When you start Democracy the first time, you'll be presented with a "How to Get Started" guide. After the first time, the player opens up to its on-line channel guide (Figure 1). Randomly selected channels from different categories will appear on the page. If the brief description under the image isn't enough, you can click the "more" link to find more information. If the show sounds interesting, subscribing to that channel is as easy as clicking the green Add button.

The sidebar to the right of the channel guide provides you with alternative ways of navigating the guide. You can search channels by keyword (for example Linux or open source) or browse the channels in different ways. Sort it alphabetically, check out the various categories, choose categories based on tags, or bring up the list of the most popular channels.

The interface is extremely easy to work with, and nothing takes away from the experience of watching the various channels. Down on the left-hand side is a small menu at the top, with a list of subscribed channels below. Democracy does start up with a handful of subscribed channels, none of which you have to keep, but these channels are a good way to get the feel of the program. If you are like me, you'll want to watch something right away, so let's do that. Click on any channel, and you get a list of the currently available episodes for that show in the main window to the left (Figure 2). This main window, by the way, is also the window where you watch your shows.

Each video is listed on the page with a small thumbnail to the left and a description of the video to the right. A link back to the site from which the video originated is often included with the description. Now, look at the thumbnail image, and you'll see a little blue down arrow in the lower-right corner. Click that arrow, and the video downloads into your collection. Remember, this isn't a streaming video application, but a digital video recorder. As the video downloads, a little progress meter appears to the left of the entry (Figure 3).

As you can see from the image, the download speed and the time remaining are both displayed along with the graphical progress bar. There's no need to wait for that video to finish downloading. You can add as many as you want, and it will all happen in the background. As you download more and more of these



videos, they will build up in your collection, which you access by clicking the button near the top in the left-hand sidebar.

When you first start the package, Democracy creates a .democracy folder in your home directory. In that folder, another folder is created called Movies. As you can probably appreciate, downloading and storing a whole lot of movies does tend to chew up your disk space over time. Depending on how your disk is organized, you may want to store your movies in another

Figure 2. Selecting a channel brings up a list of the episodes currently available for download.

Don't Gamble with Other Service Providers...Host with Hurricane Electric and Get **Fast, Reliable** Service!

**Flat Rate
Gigabit Ethernet**

1,000 Mbps of IP

\$13,000/month*

**Full 100 Mbps
Port**

Full Duplex

\$2,000/month

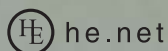
**Colocation Full
Cabinet**

Holds up to 42 1U
servers

\$400/month

Order Today!

email sales@he.net or call 510.580.4190



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

The program lets you find, download, record, manage and watch Internet television programs.

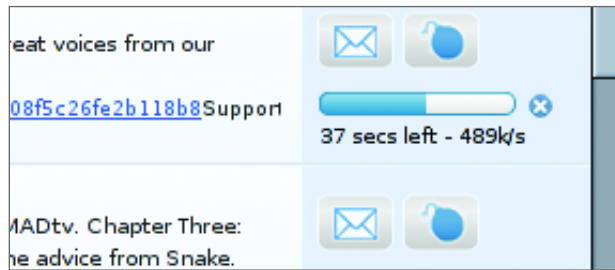


Figure 3. As shows are downloaded, a graphical meter keeps track of the progress.

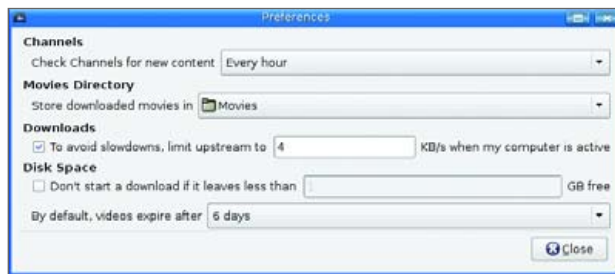


Figure 4. Most configuration options deal with space and time (storage space and expiration dates).



Figure 5. I guess you could call this "Democracy in action".

folder or even on another drive. To change the storage location, click Video on the menu bar and select Preferences. In a second or two, the Preferences dialog appears (Figure 4).

Democracy regularly checks your subscription list to check for new content. If you want, you can override the one-hour interval using the combo box at the top. What I really want to highlight, however, is the issue of storage. Click the combo box labeled Movies Directory, and you can use the file navigator to select an alternate storage location. Don't click Close quite yet though. There are other interesting settings here. Closely related to where you store your files is the amount of space you are willing to let them occupy. If you don't do anything else, two things will happen. The first is that Democracy will download and store as many movies as you desire until you run out of space. The second is that on the sixth day after you download, Democracy will clean up after itself automatically. If you would rather make that three days, go ahead and change it. The final thing I want to point out here is

the Disk Space check box (currently unchecked). If you click the box, you can specify how much space Democracy will make sure you have before it starts to download a file. When you've made all the changes you want to make, click Close to banish the dialog.

All this is about watching Internet television, so let's get back to that. After your downloads are done, click the My Collection button to see a list of all the shows that have been downloaded onto your system. Each thumbnail in the list has a little play button, a green circle with an arrow in it, that you click to play. The movie or clip will then start, using your embedded movie player (Figure 5). A slider moves along the bottom of the main window to let you know your position on the video. There's also a volume slider below the position slider.

Democracy has a nice, large viewing window, but unless you have reason not to, why not go for the big picture? To switch to full-screen viewing mode, click the full-screen button at the bottom of the screen. You'll find it to the right of the Play/Pause button.

As you watch more and more shows, you are likely going to find some that you like, others that you don't, and some that you would love to recommend to other people. Democracy gives you an option for dealing with all three of these possibilities. First, if you love a show so much that you don't want it deleted after six days, click the Save button to the left of the description (Figure 6). If you don't ever need to see this video again, click the trashcan icon. Finally, if you really, really enjoyed a particular video, and you need to tell the world, you have two choices. Click the envelope icon to send an e-mail message to your friend (or friends). The final option is to bomb it by clicking the bomb icon to the right of the video in your playlist (again, see Figure 6).

To bomb a video is actually a very good thing, and here's why. Doing this fires up your default browser (such as Konqueror or Firefox) and opens you up to the page for that video on the Videobomb.com Web site. There, your vote (bomb) will be added to others, raising the profile of that video in the list. To do this, you first need to create an account on the Videobomb.com site. Not only will this allow you to vote for the videos you enjoy, but it also provides a page to which you can direct your friends, so you can chat about the shows you bombed.

The young lady at table 12 suggests that this is yet another social networking site and she is correct. What makes this one different, however, is its tie-in to the on-demand Democracy player. In that respect, it brings the whole channel surfing, recording, watching and talking about it experience full circle. It's like hanging around the water cooler at work, chatting about what you watched last night, but in an instant gratification kind of way.

Mon Dieu! Is it that time already? Once again, the clock is telling us that closing time has arrived, *mes amis*. Of course, it won't be the first time that any of us has spent hours in front of the television until the late hours. We just don't often have this kind of selection. Perhaps François would be so kind as to refill our glasses a final time, so that we can raise a toast. Please raise your glasses, *mes amis*, and let us all drink to one another's health. *A votre santé! Bon appétit!* ■

Resources for this article: www.linuxjournal.com/article/9259.

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




Figure 6. Share videos you like with your friends by "bombing" them.



One

PGI Unified Binary™

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



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



DAVE TAYLOR

As I have said before, awk has lots of power for those people willing to spend a little time learning its ins and outs.

Analyzing Log Files Redux

If you want an easy way to calculate the amount of data transferred from a log file, you can always look awk-ward.

Last month, we spent a lot of time digging around in the Apache log files to see how you can use basic Linux commands to ascertain some basic statistics about your Web site.

You'll recall that even simple commands, such as `head`, `tail` and `wc` can help you figure out things like hits per hour and, coupled with some judicious uses of `grep`, can show you how many graphics you sent, which HTML files were most popular and so on.

More important, utilizing `awk` at its most rudimentary made it easy to cut out a specific column of information and see that different fields of a standard Apache log file entry have different values. This month, I dig further into the log files and explore how you can utilize more sophisticated scripting to ascertain total bytes transferred for a given time unit.

How Much Data Have You Transferred?

Many ISPs have a maximum allocation for your monthly bandwidth, so it's important to be able to figure out how much data you've sent. Let's examine a single log file entry to see where the bytes-sent field is found:

```
72.82.44.66 - - [11/Jul/2006:22:15:14 -0600] "GET
➔/individual-entry-javascript.js HTTP/1.1" 200 2374
➔"http://www.askdaveytaylor.com/
➔sync_motorola_razr_v3c_with_windows_xp_via_bluetooth.html"
➔"Mozilla/4.0 (compatible; MSIE 6.0;
➔Windows NT 5.1; SV1; .NET CLR 1.1.4322; .NET CLR
➔2.0.50727)"
```

There are a lot of different fields of data here, but the one we want is field #10, which in this instance is 2374. Double-check on the filesystem, and you'll find out that this is the size, in bytes, of the file sent, whether it be a graphic, HTML file or, as in this case, a JavaScript include file.

Therefore, if you can extract all the field #10 values in the log file and summarize them, you can figure out total bytes transferred. Extracting the field is easy; adding it all up is trickier, however:

```
$ awk '{ print $10 }' access_log
```

That gets us all the transfer sizes, and we can use `awk`'s capabilities to make summarizing a single-line command too:

```
$ awk '{ sum += $10 } END { print sum }' access_log
```

As I have said before, `awk` has lots of power for those people willing to spend a little time learning its ins and outs. Notice a lazy shortcut here: I'm not initializing the variable `sum`, just exploiting the fact that variables, when first allocated in `awk`, are set to zero. Not all scripting languages offer this shortcut!

Anyway, run this little one-liner on an access log, and you

can see the total number of bytes transferred: 354406825. I can divide that out by 1024 to figure out kilobytes, megabytes and so on, but that's not useful information until we can figure out one more thing: what length of time is this covering?

We can calculate elapsed time by looking at the first and last lines of the log file and calculating the difference, or we simply can use `grep` to pull one day's worth of data out of the log file and then multiply the result by 30 to get a running average monthly transfer rate.

Look back at the log file entry; the date is formatted like so: - [11/Jul/2006:22:15:14 -0600]. Ignore everything other than the fact that the date format is DD/MMM/YYYY.

I'll test it with 08/Aug/2006 to pull out just that one day's worth of log entries and then feed it into the `awk` script:

```
$ grep "08/Aug/2006" access_log | awk '{ sum += $10 }
➔END { print sum }'
78233022
```

Just a very rough estimate: 78MB. Multiply that by 30 and we'll get 2.3GB for that Web site's monthly data transfer rate.

Turning This into a Shell Script

Now, let's turn this into an actual shell script. What I'd like to do is pull out the previous day's data from the log file and automatically multiply it by 30, so any time the command is run, we can get a rough idea of the monthly data transfer rate.

The first step is to do some date math. I am going to make the rash assumption that you have GNU date on your system, which allows date math. If not, well, that's beyond the scope of this piece, though I do talk about it in my book *Wicked Cool Shell Scripts* (www.intuitive.com/wicked).

GNU date lets you back up arbitrary time units by using the `-v` option, with modifiers. To back up a day, use `-v-1d`. For example:

```
$ date
Wed Aug 9 01:00:00 GMT 2006
$ date -v-1d
Tue Aug 8 01:00:47 GMT 2006
```

The other neat trick the `date` command can do is to print its output in whatever format you need, using the many, many options detailed in the `strftime(3)` man page. To get DD/MMM/YYYY, we add a format string:

```
$ date -v-1d +%d/%b/%Y
08/Aug/2006
```

Now, let's start pulling the script together. The first step in the script is to create this date string so we can use it for the `grep` call, then go ahead and extract and summarize the bytes transferred that day. Next, we can use those values to calculate

Big-business performance for your small business.

Choose the Dual-Core Intel® Xeon® Processor in your ZT DC X9000 series system and experience outstanding performance. The better to run more applications, support more users and get more done in less time.



Hotswap Power Supplies and Fans Available.
3-Year On-site Service add \$119.00



ZT 2U Server DC X9101



ZT Tower Server DC X9102



ZT Tower Server DC X9103

- 2 x Dual-Core Intel® Xeon® Processors 5060**
(2x2MB L2 Cache, 3.20 Ghz, 1066MHz FSB, Intel® HT, EM64T)
- Windows® Server 2003 Standard R2b w/ 5 CAL add \$699.00
 - Intel® 5000V Chipset Server Board
 - Support up to 16GB DDRII 533 FB-DIMM Memory
 - SAS / SATAII Max. Storage 4.5TB
 - 6 x 1" Hot-swap SAS / SATAII Drive Bays
 - Slim DVD-ROM and Floppy Drive
 - 4 x SATAII Ports via ESB2 Controller (RAID 0, 1, 5, 10 support)
 - Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
 - 2U Rackmount Chassis w/ 550Watt Power Supply
 - SuperDoctor III Server Management Software
 - 3-Year Parts and Labor Limited Warranty
 - 3-Year On-site Service add \$119.00

- Dual Core Intel® Xeon® Processor 5030**
(2x2MB L2 Cache, 2.66 Ghz, 667MHz FSB, Intel® EM64T)
- Windows® Server 2003 Standard w/ 5 CAL ADD \$699.00
 - Intel® 5000V Chipset Server Board
 - Support up to 16GB DDRII 533 FB-DIMM Memory
 - SAS / SATAII Max. Storage 3TB
 - 4 x 1" Hot-swap SAS / SATAII Drive Bays
 - 16x Dual-layer DVD-RW and Floppy Drive
 - 4 x SATAII Ports via ESB2 Controller (RAID 0, 1, 5, 10 support)
 - Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
 - Mid Tower Server Chassis w/ 645W Heavy Duty Power Supply
 - SuperDoctor III Server Management Software
 - 3-Year Parts and Labor Limited Warranty
 - 3-Year On-site Service add \$119.00

- Dual Core Intel® Xeon® Processor 5140**
(4MB L2 Cache, 2.33 Ghz, 1066MHz FSB, Intel® HT, EM64T)
- Windows® Server 2003 Standard w/ 5 CAL ADD \$699.00
 - Intel® 5000V Chipset Server Board
 - Support up to 16GB DDRII 533 FB-DIMM Memory
 - SAS / SATAII Max. Storage 3TB
 - 8 x 1" Hot-swap SAS / SATAII Drive Bays
 - 16x Dual-layer DVD-RW and Floppy Drive
 - 4 x SATAII Ports via ESB2 Controller (RAID 0, 1, 5, 10 support)
 - Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
 - 4U Rackmount / Tower Convertable Chassis w/ 645Watt Redundant-Cooling Power Supply
 - SuperDoctor III Server Management Software
 - 3-Year Parts and Labor Limited Warranty
 - 3-Year On-site Service add \$119.00

Starts at **\$1,799**

Starts at **\$999**

Starts at **\$1,599**

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

Go to **ztgroup.com/go/linuxjournal**
Call **866-ZTGROUP (866-984-7687)**

Promote code : Lj1106



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

The other neat trick the date command can do is to print its output in whatever format you need, using the many, many options detailed in the `strftime(3)` man page.

other values with the `expr` command, saving everything in variables so we can have some readable output at the end. Here's my script, with just a little bit of fancy footwork:

```
#!/bin/sh

LOGFILE="/home/limbo1/logs/intuitive/access_log"

yesterday="$(date -v-1d +%d/%b/%Y)"

# total number of "hits" and "bytes" yesterday:

hits="$(grep "$yesterday" $LOGFILE | wc -l)"

bytes="$(grep "$yesterday" $LOGFILE | awk '{ sum += $10 } END { print sum }')"

# now let's play with the data just a bit

avgbytes="$(expr $bytes / $hits )"
monthbytes="$(expr $bytes \* 30 )"

# calculated, let's now display the results:

echo "Calculating transfer data for $yesterday"
```

```
echo "Sent $bytes bytes of data across $hits hits"
echo "For an average of $avgbytes bytes/hit"
echo "Estimated monthly transfer rate: $monthbytes"
```

```
exit 0
```

Run the script, and here's the kind of data you'll get (once you point the `LOGFILE` variable to your own log):

```
$ ./transferred.sh
Calculating transfer data for 08/Aug/2006
Sent 78233022 bytes of data across 15093 hits
For an average of 5183 bytes/hit
Estimated monthly transfer rate: 2346990660
```

We've run out of space this month, but next month, we'll go back to this script and add some code to have the transfer rates displayed in megabytes or, if that's still too big, gigabytes. After all, an estimated monthly transfer rate of 2346990660 is a value that only a true geek could love! ■

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.

ISPCON

THE INTERNET INDUSTRY EVENT

WHAT'S HOT IN LINUX



See all the hottest **gear**, coolest **software** and latest **innovations** that will *turn your business up a notch!*



Get **solutions** to your key challenges first-hand from industry hotshots and your peers who are deploying the latest strategies.

Discover it first at ISPCON Fall 2006

Register today with **Customer Code LNXF6**
to get your **FREE EXHIBIT HALL PASS!**

NOVEMBER 7-9, 2006 • SANTA CLARA CONVENTION CENTER • WWW.ISPCON.COM

SUPERMICRO[®]

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

SuperServer 6015X-8/T

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



SUPER[®] Serverboard X7DBX-8/T

Supermicro's Feature Advantages

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



Dual-Core Xeon 5100/5000 Series

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

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

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

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

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

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

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

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



MICK BAUER

Whereas chrooting restricts a process to a subset of the host system's real filesystem, virtualization restricts the process to a complete virtual machine running within the host (real) machine.

Running Network Services under User-Mode Linux, Part I

Leverage the Linux kernel's virtualization features to isolate network daemons.

In my May 2006 Paranoid Penguin column, I expounded on the virtues of Debian 3.1's excellent support for virtualization environments, including User-Mode Linux. In that same issue, in the article "User-Mode Linux", Matthew Hoskins gave a quick-and-dirty recipe for test-driving User-Mode Linux using prebuilt UML kernels and root filesystem images.

Did those articles whet your appetite for a more comprehensive and security-focused look at using UML? If so, you're in luck; for the next couple of columns, we're going to dive into the User-Mode Linux experience and cover every step (including every command) for creating your very own User-Mode Linux containers for network services.

Objectives

So, why are we doing this, and what do we hope to achieve? As I've said before in this space, virtualization is similar to the concept of the chroot (changed root) jail. It encapsulates a process or daemon into a subset of the operating environment in which it resides, in a manner that makes it much harder for attackers to get at the underlying environment should they succeed in compromising that process or daemon.

Whereas chrooting restricts a process to a subset of the host system's real filesystem, virtualization restricts the process to a complete virtual machine running within the host (real) machine. This includes a completely virtualized hard disk, memory and kernel, and even virtualized system devices, such as network and sound cards. In the case of User-Mode Linux, this is achieved by running a guest (virtual) kernel as a user-space process within the host (real) kernel.

Because both guest and host kernel are Linux kernels, virtualization in User-Mode Linux is fast and efficient. And, because the guest kernel does *not* need to run as root under the host kernel, even attackers who compromise some daemon on the guest system *and* escalate their privileges to root (on the guest system) *and* somehow manage to gain shell access to the underlying host system will have achieved only unprivileged access to that host system.

This does not make it impossible to gain root access to the host system. If attackers do make it as far as shell access on the host, they may be able to escalate their privileges via some local privilege escalation vulnerability in the host's kernel or some user-space program on the host. (Remember: no vulnerability is strictly local on *any* networked system!) It *does* mean, however, that it's more difficult for attackers to get to the point of being able to exploit such a vulnerability, especially if it isn't also present on the guest (virtual) system.

This brings us to our design goals. The guest machine should be as bare-bones as possible with respect to

installed software—both to minimize resource utilization and to minimize its potential for compromise (its attack surface). If, for example, the guest machine is to act as a DNS server, it should have basic network support, BIND (or some other DNS server package) and very little else. No X Window System and no Apache—nothing else not directly related to DNS services.

If you're really paranoid, you even can skip the Secure Shell daemon and instead administer the system via a virtual serial console. (Though allowing SSH from only authorized IP addresses, such as that of the host system, might be a more reasonable middle ground.) You could also run User-Mode Linux under SELinux; however, that's beyond the scope of this series of articles.

If a single bastion server is to host multiple network services—for example, Apache and BIND—you could run two different guest systems on the same host: one containing only Apache and its dependent packages and another containing only BIND et al. In this way, a vulnerability in BIND would not lead directly to Web site defacement. Conversely, a poorly coded Web application would not necessarily lead to DNS tampering.

In summary, our two design principles will be to run one virtual machine per major network service and to make each virtual machine as minimal and secure as possible. The end result will (hopefully) be a very compartmentalized bastion server that places as much defensive abstraction as possible between attackers and total root compromise.

For the remainder of this series of articles, I use the example of a single guest system running BIND. Both guest and host system are based on Debian 3.1, because Debian is so popular for UML guests (it lends itself to stripped-down installations—a trait it shares with Slackware). However, most of what follows also applies to other Linux distributions on both host and guest.

Our tasks are:

1. Build a host kernel optimized for hosting User-Mode Linux guests.
2. Build one or more guest kernels to run on top of the host.
3. Obtain and customize a prebuilt root filesystem for our guests.
4. Run, configure and harden our guest system for secure DNS services.

Preparing the Host

First, you need to make sure you've got the right kind of kernel on your host system. You very likely may need to



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

AoE is your answer!

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

www.coraid.com



EtherDrive® SR1520

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



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

Visit us at www.coraid.com
for more information.



1.706.548.7200

The Linux Storage People

www.coraid.com

Thanks to the magic of COWs, it's therefore possible to run the same guest kernel and root filesystem combination multiple times, by defining a unique COW file per instance.

compile a new kernel.

On the one hand, some Linux distributions already have User-Mode Linux compiled into their default kernels. On the other hand, your distribution of choice may or may not also have the skas (separate kernel address space) patch compiled in as well. It is, in fact, somewhat unlikely that your default kernel has skas support. Although the Linux kernel source code has included UML support since version 2.6.9, the skas patch is still maintained separately (and Linus has resisted its inclusion).

The skas patch is important. It greatly improves UML performance and security by running the guest system's kernel in separate address space from its other processes (just like the host's kernel does). The User-Mode Linux Web site's skas page on SourceForge provides a more detailed explanation of why you need skas (see the on-line Resources).

To obtain kernel source code, your best bet may be simply to install your Linux distribution's kernel-source package. Take care, however, that your distribution provides a kernel version of 2.6.9 or higher, because UML support is included from 2.6.9 onward, and prior UML patches had security vulnerabilities.

Because Debian 3.1 still uses kernel version 2.6.8, I decided not to use the official Debian kernel packages and instead downloaded the 2.6.17 kernel from kernel.org. I did, however, install the kernel-package package, which

provides tools for generating Debian packages from official kernel source.

Besides kernel source code, you need the skas patch, the latest version of which is available on Blaisorblade's site (see Resources). Be sure to download the patch version that corresponds to the kernel source code you're about to patch.

On my Debian host, I unpacked my official source code to `/usr/src/linux-2.6.17.3`, renamed the source code directory to `/usr/src/linux-2.6.17.3-host` and copied the skas patch tarball (`skas-2.6.17-rc5-v9-pre9.patch.bz2`) to `/usr/src`. I then changed ownership of the directory `/usr/src/linux-2.6.17.3-host` to a nonroot account. (Adhering to the principle of never being root unless you really need to, we're going to do most of this kernel build as an unprivileged user.)

Here are the commands I executed as root:

```
host:/usr/src/# tar -xjvf ./linux-2.6.17.3.tar.bz
host:/usr/src/# mv ./linux-2.6.17.3 ./linux-2.6.17.3-host
host:/usr/src/# chown mick ./linux-2.6.17.3
host:/usr/src/# su - mick
```

To apply the skas patch, I then navigated, as my nonroot user, to `/usr/src/linux-2.6.17.3-host` and ran the following command:

Keeping Your Kernels and Guests Straight

In the contexts of User-Mode Linux, VMware and other virtualization systems, we use the words host and guest in a very specific way. Your host is the system that runs the virtualization environment—that is, it acts as a host to one or more virtual machines. Guests are virtual machine instances that live on top of the host.

Therefore, when we speak of the host kernel and guest kernels, remember that guest kernels run on top of the host kernel. In User-Mode Linux, your host kernel is a normal Linux kernel, compiled for your particular hardware platform (Intel x86, IBM PowerPC and so on), with User-Mode Linux features (including the optional skas patch) compiled in as well.

Your guest kernel, on the other hand, must be compiled to run on virtual hardware: the um architecture. Other than that, it does *not* need the skas patch or User-Mode Linux support enabled. Unless, that is, you want to run other guest kernels on top of it. Running guests within guests is possible (this is called nesting), but well beyond the scope of this article.

Each UML virtual machine instance consists of a guest kernel, a guest root filesystem and a COW (Copy On Write) file. The root filesystem is a disk image file; it contains every file in your virtual machine except the kernel itself. When you execute a guest kernel, the root filesystem file is mounted in precisely the same way you'd mount any other disk image, for example, a CD ISO file. Like a CD-ROM, it's used in read-only mode. Any changes you make to the virtual filesystem in the course of a UML session, including new files and file deletions, are stored in a COW file.

Thanks to the magic of COWs, it's therefore possible to run the same guest kernel and root filesystem combination multiple times, by defining a unique COW file per instance.


```
host:/usr/src/linux-2.6.17.3-host$ bunzip2 -c
➔ ../skas-2.6.17-rc5-v9 -pre9.patch.bz2 | patch -p1
```

Next, from the same directory, I issued the command `make menuconfig`. When setting up the kernel configuration for User-Mode Linux, the defaults generally are fine, though you should ensure that the configuration matches your host's hardware. In addition, it's probably prudent to double-check the following settings:

- Under Processor type and features, make sure `/proc/mm` is enabled.
- Under Networking options, make sure IP: tunneling and 802.1d Ethernet Bridging are enabled. If you intend to restrict guest system behavior with iptables, you also may want to check the Network packet filtering section to ensure that Core Netfilter Configuration, IP: Netfilter Configuration and Bridged IP/ARP packets filtering are set up.
- Under Network device support, enable Universal TUN/TAP device driver support.
- And, by all means, make sure to hard-compile (into the kernel, not as a module) the filesystem in which your system's root partition is formatted (for example, ext3 or ReiserFS).

From this point on, the process is the same with any other kernel build: issue the commands `make bzImage` and `make modules`; then, become root and issue the commands `make modules`, `make modules_install` and `make install`. (Or in the case of Debian, use the `make-kpkg` command to achieve the same thing, and run `dpkg` to install the resulting kernel package.)

Once your new host kernel is installed, reboot your system. Your host system is now capable of running User-Mode Linux guest systems.

Creating a Guest Kernel

Okay, we've got UML host capabilities, but we still need a guest kernel to run. This process is somewhat simpler than the host-kernel build, because we don't need the `skas` patch.

First, navigate back to the directory in which your Linux kernel-source tarball resides, and unpack it a second time. Remember when we renamed the unzipped source code directory? This was so we could unpack the kernel tarball a second time. We need to build our host and guest kernels in separate source trees.

On my Debian test system, therefore, I unpacked the source tarball to `/usr/src/linux-2.6.17.3`, and this time, renamed it to `/usr/src/linux-2.6.17.3-guest`. Again, change ownership of this directory to a nonprivileged user, and change your working directory to it.

Again, at this point we can skip the step of applying the `skas` patch. Because we're going to compile our kernel for the special `um` (User-Mode Linux) architecture rather than for a real architecture like `x86`, I recommend you prepare your source code tree with the following three commands:

```
host:/usr/src/linux-2.6.17.3-guest$ make mrproper ARCH=um
host:/usr/src/linux-2.6.17.3-guest$ make defconfig ARCH=um
```

```
host:/usr/src/linux-2.6.17.3-guest$ make menuconfig ARCH=um
```

The `make mrproper` command clears out any configuration and object files in your source tree; `make defconfig` generates a fresh default configuration file appropriate to the `um` architecture; and `make menuconfig`, of course, gives you the opportunity to fine-tune this configuration file.

Pay particular attention to the following:

- Life will be simpler if you skip loadable kernel module support and hard-compile everything into the kernel. If you really want kernel modules, see the User-Mode Linux HOWTO, Section 2.2 (see Resources).
- Under Processor type and features, double-check that your system architecture is set to `um` (User-Mode Linux), and make sure `/proc/mm` is enabled.
- Under Networking options, make sure IP: tunneling and 802.1d Ethernet Bridging are enabled.
- Under Network device support, enable Universal TUN/TAP device driver support.
- Disable as many of the specialized hardware kernel modules as possible; this kernel is going to be running on virtualized hardware, so you won't need support for wireless LAN hardware, obscure parallel-port devices and so forth.

Once you've saved your new configuration file, you can compile the kernel with this command (*without* first becoming root; execute this as an unprivileged user):

```
host:/usr/src/linux-2.6.17.3-guest$ make linux ARCH=um
```

Note that I did *not* tell you to make a zipped or bziped image. Remember, you're going to be running this kernel as though it were a user-space command, so it shouldn't be compressed. The finished kernel will be located in the top-level directory of your source tree (`/usr/src/linux-2.6.17.3-guest` in the above examples) and will be named `linux`—you'll probably want to rename it to something more descriptive, such as `uml-guestkernel-2.6.17.3`. You'll also probably want to move it to the directory from which you intend to run it—perhaps something like `/usr/local/uml/`.

By the way, don't be scared by the size of your guest kernel file. Most of that bulk is symbol information that will *not* be loaded into memory when you execute it.

Conclusion

Your host system now fully supports User-Mode Linux, and you've got a guest kernel image to run. The next step is to obtain or create a root filesystem image to use with the guest kernel. That's where we'll pick up again next time! ■

Resources for this article: www.linuxjournal.com/article/9260.

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".



JON "MADDOG" HALL

A Small Conference

It takes creative thinking, focus and, most of all, plenty of lead time to create a successful free software event.

"What are you working on?", I asked Dennis, a young friend of mine from Florianópolis. "I am working on the design of a T-shirt for the conference we are developing, but I do not have any good ideas for the design", he said. "What is the purpose of the conference?", I asked. "Who is the target market?" "What do you mean?", he asked suspiciously, probably thinking that this was going to turn into a lecture on marketing. He was right.

Many Linux User Groups (LUGs) have tried to put on small, local events to introduce people to Free and Open-Source Software (FOSS). Some have been successful, and some have failed. Some have even been "too successful", over time burning out the volunteer staff that put the conference together. But, the successful ones always have tended to follow a similar pattern—that of thorough planning.

"Defining your purpose and who you are going to be reaching is very important", I said, and continued:

For example, do you want the conference to be technically oriented, to satisfy programmers or systems administrators, or do you want it to be more business-oriented, to convince business people that they should be using free software? Is your aim to show people how they can create jobs or make money with free software? You can do all of this with one event, but it would be a larger event and much more difficult to do than to concentrate on just one audience.

Dennis thought about this for a while, and said, "I want it to be technical, but invite a few business people." "And when do you want to have this conference?", I asked. "In two months", he said.

One of the biggest mistakes a group makes in planning an event is trying to have it too soon. Often I get an invitation to speak at an event three months away. I tell the people that I would have been willing to attend, but that the date has been booked for six months. Many venues big enough to hold even a small conference are often booked six to nine months in advance. To have the most leeway, you probably should start planning a year in advance. It will not be constant planning during that year, but the bigger items (venue, keynote speakers and so forth) should take precedence early in the planning process.

"What is the theme of your event?", I asked. "Free Software" came the reply.

In the past, "Free Software" meant "Linux" or "BSD" or some of the GNU tools, but today, "Free Software" means audio/video tools, customer relationship management software, Voice over IP, content management systems, TV capture and playback, development tools, many types of database programs, clustering software and much, much more. Trying to cover such a wide range of interests is difficult in a small conference.

"Why not focus on just 'audio/video'?", I suggested. "You could even invite Gilberto Gil [Brazil's most famous rock star, and the Minister of Culture under President Lula] to speak." This went over well with the rest of the planning committee, who was now beginning to gather around the T-shirt table.

Dennis had taken steps to avoid the second big mistake that a lot of groups make—a planning committee that is too small. "Many hands lessens the load", and in a lot of ways make planning the event much more fun. By enlisting a set of enthusiastic friends, Dennis had committees for each of the major functions that they needed. A quick status recount found that any suitable venue was only available nine months in the future. The "two-month wonder conference" would have to be rescheduled for later, which was greeted with a sigh of relief from the program committee of Chico, Douglas and Felipe.

"We have been having trouble finding speakers", they wailed. "Even with the new dates, we do not know who to ask." I suggested that with the new dates and the new theme of "audio/video" that they might approach the developers of some





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



of the projects to see if they would be willing to come. If, however, their target audience is users rather than developers, they would be better off getting local people to learn the projects, and then do a presentation on how they work and how to use them rather than how to develop them. Other good topics, particularly for those new to Free Software, are the general topics of how Free Software works, the different software licensing models and how to turn in a good bug report. Using local people to do a lot of the presentations keeps down the costs of hotel rooms, airline fares and translators (if speakers do not speak your language), and it also gives local people a chance to develop speaking skills—useful the next time you put on a conference.

“But without big-name speakers, how do we attract sponsors? We do not want to charge any money for our event”, said Rodolfo, the treasurer.

“Although I appreciate that you want everyone to come to the event ‘for free’, there are times when you need to charge a little bit to cover costs”, I said, and added:

Why not ask for R\$5 (five Reals) as a donation? You may be surprised how much money you get if people have a good time and learn a lot. Sponsorships also can come in the form of Internet connectivity, equipment loans, advertising and other trades, which are easier to get than money. Also, local events can attract donations and vendor sales from local vendors. A local bookstore, for instance, can stock up on Free Software books if they are given enough warning and a list of books people might want to buy. And remember, that although many people think of large companies when they think of sponsorships, small companies also can contribute smaller sums of money that can go a long way

toward paying the costs of a great conference.

“I went to a conference that was free, but they sold a T-shirt and a CD-ROM of the conference proceedings for a ‘donation’ of R\$10”, said André. “It could not have cost them more than R\$5 to make both, so each person contributed R\$5 to the conference.”

“And, raffles of donated products as prizes is also a way to make money. I remember a conference that made 3,000 Australian dollars off raffled prizes that were donated by vendors”, I said, and continued:

The main thing, however, is to try to keep the costs down. Because most conferences are educational, a lot of times you can get the local college or university to donate the space, with the cost of custodial and security personnel as the only charge. These days, a lot of such institutions are happy to help you plan a Free Software event. They see it as good for their students and faculty and also for their image with potential students if done well.

“What about food?”, asked Henry. “You need food for people to eat.” “A picnic basket or a brown-bag lunch is fine”, I said, adding:

A cooler of sodas and bottled water sold at a reasonable price, and you will have a lot of happy people. Make sure you have trash cans around for the wrappings, and recycle the cans and bottles. One of the best conferences I have attended sold a simple loaf of French bread with a little spread on it for one Euro, and a bottle of drink for another Euro. It was enough for lunch, and fast to eat. And, of course, you’ll need coffee.

“Hotel rooms”, said André, “we will need them. How do we handle the hotel rooms?” I answered:

For a one-day event, with local attendees, not too many hotel rooms will be needed. Most people can make their own reservations simply by having a list of reasonably priced hotels local to the event. For guest speakers, you may want to make the hotel reservations for them to help keep costs under control, but many guest speakers in the Free Software world are happy to stay in a host’s house, or a university dorm room, or some other such place that is clean and quiet with high-speed Internet available.

André smiled, because I had just described his home where I had stayed during a Software Livre conference.

“What about the business people?”, asked Dennis. I agreed that the business people needed a slightly different approach, but that we should talk about that tomorrow, as it was getting late, and the T-shirt design was still not finished. “What about a Tux riding a surfboard?”, I asked, “or maybe in a beach chair with a chimarrão in its flipper?” The rest of the conference committee crowded around while we sketched examples.

Aside: one of my favorite times of the year is “OpenBeach”. This is a small get-together that is usually created for a group of Free Software people to discuss things that are happening, but also to enjoy each other’s company and meet the families (spouses, children and so on) of the people they may deal with only by e-mail and chat on a day-to-day basis. It is an event by the seaside, with wireless Internet abounding. This year, it is preceded by an event at the Universidade Federal de Santa Catarina on December 6th, 7th and 8th, with a natural follow-over to OpenBeach. The combination of the conference and the relaxing weekend following it will be fun—and when putting together an event, *fun* is one of the most important things. ■

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

Need help,
not headaches?

OffAssist

All of your business needs. One easy solution...

BOOKKEEPING
and
VIRTUAL ASSISTANCE

Get your FREE copy of “VA’s and the Tech Sector: Plugging a VA into Your Business” at:
www.offassist.com/linuxjournal

www.offassist.com

We Speak Geek! Open Source — Open Mind

EXPAND YOUR HIGH PERFORMANCE COMPUTING FLEXIBILITY

ENJOY GREAT DUAL-CORE PERFORMANCE PLUS FLEXIBILITY TO RUN 32-BIT AND 64-BIT APPLICATIONS WITH DUAL-CORE XEON PROCESSORS IN SERVERSDIRECT SYSTEM.



STARTING PRICE
\$2,499

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

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

FEATURE HIGHLIGHTS:

Intel Xeon 5050 3.0GHz Dual Core Processor (Dual Processor Option)
 3U Chassis with 800W Redundant Power Supply
 Intel® 5000P (Blackford) Chipset
 Kingston 1024MB 667MHz DDR2 ECC FB-DIMM(2pcs x 512MB)
 Seagate 400GB SATA-II 16MB Cache 7200RPM

16 x 1" Hot-swap SATA Drive Bays
 ATI ES1000 Graphics with 16MB video memory
 Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
 RAID 0, 1, 5, 10 Support

1U ENTRY LEVEL SERVER

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



1U Rackmount Chassis with 300W Power Supply
 Intel® Pentium® D 930 3.0GHz Dual Core 800FSB
 Intel® E7230 (Mukilteo) Chipset
 Kingston 1024MB 667MHz DDR2 ECC (2pcs x 512MB)
 Seagate 250GB SATA-II 7200RPM hard drive
 2 x 1" Hot-swap SATA Drive Bays
 ATI RageXL graphics, 8MB
 2x Intel® 82573V (Tekoa) PCI-e Gigabit LAN Ports
 RAID 0, 1

\$999

SKU# SDR-1110T

COST EFFECTIVE 2U SERVER

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



Intel Xeon 5050 3.0GHz Dual Core Processor (Dual Processor Option)
 2U Chassis with 600W Power Supply
 Intel® 5000V (Blackford VS) Chipset
 Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2pcs x 512MB)
 Seagate 400GB SATA-II 16MB Cache 7200RPM
 6 x 1" Hot-swap SATA-II Drive Bays
 ATI ES1000 Graphics with 16MB video memory
 Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
 RAID 0, 1, 5, 10 support

\$1,599

SKU# SDR-2501T

4U TOWER/RACKMOUNTABLE 64-BIT XEON ENTERPRISE SERVER

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



Intel Xeon 5050 3.0GHz Dual Core Processor (Dual Processor Option)
 4U Rackmountable / Tower
 Intel® 5000X (Greencreek) Chipset
 Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2pcs x 512MB)
 8pcs x Seagate 400GB SATA-II 16MB Cache 7200RPM
 3Ware 9550SX-8port RAID Controller Card
 ATI RageXL graphics, 8MB
 GeForce 6200 TC 128MB OB 256M support DDRTV-out PCI-E
 Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller

\$3,999

SKU# SDP-7045A-TB

SDR-5500T 5U ADVANCED STORAGE SERVER

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



Intel Xeon 5050 3.0GHz Dual Core Processor (Dual Processor Option)
 5U Chassis with 24 hot-swap bays & 950W redundant power supply
 Intel® 5000P (Blackford) Chipset
 Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2pcs x 512MB)
 12pcs x Western Digital 400GB SATA RAID Drive
 1pc x 3Ware 9550SX-12 port RAID Controller Card
 ATI ES1000 Graphics with 16MB video memory
 Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller
 RAID 0, 1, 5, 10 support

\$6,799

SKU# SDR-5500T

SERVERS DIRECT CAN HELP YOU CONFIGURE YOUR NEXT5 HIGH PERFORMANCE SERVER SYSTEM - CALL US TODAY!

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

1.877.727.7887 | www.serversdirect.com





DOC SEARLS

The Search for Terrestrial Stupidity

How about a SETI project to build knowledge about how the Net is actually working?

Tom Evslin wants us to step back from the Net Neutrality fray and resolve the issues through widespread knowledge that currently does not exist.

If Net Neutrality is a good thing, shouldn't we be able to test for it? Shouldn't everybody on the Net be in a position to see how things are going for them? And, wouldn't it be useful to slice and dice data coming in from all those nodes looking at Net performance from the edges?

Those were some of the questions raised by Tom Evslin with "Net Neutrality at Home: Distributed Citizen Journalism against Net Discrimination"—a recent luncheon talk at Harvard's Berkman Center for Internet and Society (see the on-line Resources). "The goal of what I'm proposing is to preserve in the United States an Internet that is equally open to all applications regardless of who owns the network and regardless of who the application owner is." He adds, "I'm being US-centric here because we suck as far as the rest of the world is concerned....The problem is here." He also adds, "Note that this doesn't mean that all applications will function equally well on every network."

As an example, he gives his own Internet connection from rural Vermont, which bounces off a satellite 25,000 miles over the equator and involves unearthly latencies that make it nearly unsuitable for VoIP. (Ironically, Tom is a VoIP pioneer. At the time he sold his wholesale VoIP company several years ago, it was the #7 carrier of voice data traffic minutes in the world.) A neutral network would make a best effort to deliver packets without discrimination in favor or against its source, destination or content. As Tom puts it, "What we want to see is each network equally open to applications, and not be more open to the application of the network owner, particularly if the network owner happens to be a monopoly."

This is where the line between technology and politics blurs. Carriers and other neutrality opponents say the Net itself has never been neutral and has always allowed many kinds of discrimination. They argue that some applications—live teleconferencing, VoIP, streaming audio and video, fault-tolerant grid computing and live remote surgery, for example—would all benefit from QoS (Quality of Service) efforts that are anything but "neutral". And they point out that discrimination of all sorts—in provisioning asymmetries, multiple service levels, selective port blockages and specific usage restrictions, to name a few—have been common practices for nearly as long as ISPs have been in business. They'd like to retain the right to discriminate, or to improve service any way they please, and to charge customers willing to pay for the benefits. They say they'd like to do that without government interference (even though carriers inhabit what they call "the regulatory environment").

Meanwhile, neutrality advocates, such as Web inventor Tim Berners-Lee, want laws to preserve the neutrality they say has always been there and is threatened by carriers who loathe the concept. Plus, it's obvious (except to those employed by the carriers—a population that sadly includes many lawmakers) that the carriers have little if any interest in building open infrastructure that enlarges business opportunity for everybody who builds on it. They would, in every case, rather capture markets

than liberate them—even if they would clearly have privileged first-mover and incumbent positions in those liberated markets. To them, "free market" means "your choice of silo".

Tom Evslin wants us to step back from the Net Neutrality fray and resolve the issues through widespread knowledge that currently does not exist. Specifically, he'd like as many users as possible to test their network connections for upload and download speeds, DNS speed, latency, jitter, blocking, consistency and uptime, to name a few of many possible variables.

Yes, techies can run some of these tests at the command line (with ping, traceroute and so on). And today, any user can visit a site such as Speakeasy.net or BroadbandReports.com to test upload and download speeds in a browser. (BroadbandReports even lets users compare results with those of other customers of the same provider.) But Tom wants to go much further than that. He wants everybody to know what they're getting and to pool data that will paint clear pictures of how individual networks and network connections are performing over time. He believes this will not only provide useful information to both sides of the current debate, but will allow everybody to observe and speak about the Internet with far more understanding than individuals have today.

"We don't want to look just for discrimination", Tom says. "We want the result of running the tools to be sort of a consumers' report map of Internet quality in general....The tools can measure both quality, and then discrimination as an aspect of quality—if the discrimination exists. But even if there's no discrimination, we'll get useful data over what kind of quality to expect where." He sees much to gain and little to lose for everybody. That is, if everybody—or at least a very large number of users—participates.

A number of questions then follow:

1. Exactly what kind of tests are we talking about?
2. How do we get users to participate on a large or massive scale?
3. If millions of users are running millions of tests or probes, how do we prevent what we might call an "insistence on service attack"?
4. How do we compile, edit and publish results?

One answer to the first question came from a report about Dan Kaminsky releasing details about a traceroute-like TCP-based fault probe, at the Black Hat security conference in August 2006. The report says:

But unlike Traceroute, Kaminsky's software will be able to make traffic appear as if it is coming from a particular carrier or is being used for a certain type of application, like VoIP. It

will also be able to identify where the traffic is being dropped and could ultimately be used to finger service providers that are treating some network traffic as second-class.

Look for this capability amidst a free suite of tools called Paketto Kieretsu Version 3. Now, what else?

For guidance, Tom says the tools must be:

- Verified and calibrated.
- Open source.
- Perceived as safe.
- Do non-destructive testing.
- Return value to each user.

One model he brought up was SETI@home. Here thousands of individuals contribute otherwise idle compute cycles to the Search for Extraterrestrial Intelligence (SETI) Project. That's a familiar model to many of us, but not likely to attract users who aren't turned on by the challenge of helping find ET. So Tom is looking for something that is SETI-like in distribution, but pays off with practical information for

the users. The following are some questions from my notes at the luncheon:

- What if users actually knew how well the Net and its providers worked for them, on both absolute and relative scales?
- What if users could look at their connection speeds the same way they look at speedometers in their cars? (Speakeasy.net does something like this with its speed tests, but how about making the test independent of any company?)
- What if users could monitor packet loss or link quality with the same ease as they check signal strength on a cell phone?
- What if users could see by a simple indicator that the Wi-Fi connection at the conference they're attending won't allow outbound e-mail? (How about a list of port blockages and what they mean?)

The program would have to be widely distributed. Tom says:

We want volunteers to run servers, to make sure various ports are open and to test the geography—like for DNS

Expert Included.

Chris brings years of experience with servers and server components to developing innovative, reliable solutions. He selects the best components, tests compatibility, and validates configurations to create solutions that exceed the expectations of our diverse customers.

Chris likes the Rackform iServ R266 with two Dual-Core Intel® Xeon® 5100 Series Processors because he knows that its CPUs provide better performance while using less power and running cooler. The iServ R266 meets his uncompromising standards: fast, efficient, expandable, and shiny.

When you partner with Silicon Mechanics, you get more than a finely tuned Intel solution – you get an expert like Chris.



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

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

The US has been falling behind the rest of the civilized world in broadband speed and penetration.

propagation. We need people who are willing to have their servers be the proxy for testing the intentional degrading of file sharing, SIP, P2P protocols and geography. Because geography is an issue. Countries now have firewalls. There might be legitimate peering problems, or routing issues. But we need to know when actual blocking is going on.

Where would these tools come from? The obvious answer is the Free Software and Open Source communities. "It is absolutely essential that the tools we get be open source", Tom says. "The tools themselves might be prejudiced. So you need to be able to see inside them to know that they're not. Second, we want to be able to bring to bear as much of the technical community as cares to participate in the development and elaboration of these tools." Tom thinks the application vendors should contribute to the effort as well, because they could only benefit from knowledge about the network. Same goes for the carriers, who would presumably like to gain bragging rights about how well they perform.

There needs to be organizations, perhaps on the SETI model, "so the task of information collection and analysis is distributed, as well as just the initial probing", Tom says. Also:

We need people responsible for verification....I'm very sensitive to that, because I've been wondering whether my satellite ISP is blocking Skype. I go on Skype and the BlueSky forums and see one person saying, "I ran this test that shows absolutely that there's been blocking", and another person saying, "The application is working for me but the test is failing"....So it's not a simple thing to know a test is actually working. One particular problem with Skype, and why Skype might not benefit from this as well as other providers, is that Skype uses a proprietary protocol....It's hard to imagine Skype contributing the particular open-source tool that is necessary to debug the things that might happen to the protocol that they're keeping secret.

Then again, having these kinds of tools looking at the network would help expose to users the deficiencies, in an open world, of closed protocols, codecs and other techniques for maintaining silos and keeping customers captive.

David Isenberg points out, "Unless you have tools for each application, that are app spec, you always run the risk that the test works fine in a generic sense and then they've got this deep packet inspection that finds the signature of the given application and blocks it." Tom answers, "So you'd like to have a tool where you could feed in the signature of the application and test that generically, and at the same time you'd like to test the protocols that they use. SIP makes sense as an example."

There is an editorial function too. News needs to go out through traditional media, as well as bloggers and other Net-based writers. The end result, in addition to keeping the carriers honest, is a far more well-informed public. Right now, most users know far less about how they travel the Net than they do about how they travel the road system. "Latency", "jitter", "packet loss" and "port blocking" are no more technical than "speed", "acceleration", "stopping distance" or "falling rock zone". Network performance knowledge should be common, not professionally specialized.

The US has been falling behind the rest of the civilized world in broadband speed and penetration. Japan and Korea are committed to making fiber-grade service available to their entire populations, and other countries are similarly motivated

to do what the US still cannot, because most of its Internet service is provided by a duopoly that cares far less about providing Net infrastructure than about delivering high-definition TV. Clearly, no help will come from lawmakers who still think a highly regulated phone/cable duopoly is actually a "free market" for anything, much less the Internet.

David Isenberg wrote his landmark paper, "The Rise of the Stupid Network" (see Resources) when he was still working for (the original) AT&T. The paper observed that most of a network's value is on its edges, rather than in its middle. At the time AT&T was busy engineering intelligence into its switches and other mediating technologies. Meanwhile, Dr Isenberg said that the network should be stupid (say, in the same way that the core and mantle of the Earth is stupid). It should be there to support the intelligence that resides on it and takes advantage of it, but is not reducible to it. In 1998, when he wrote that essay, the Net was already well-established. Yet the thinking of the carriers was still deeply mired in the past. Here's the gist of the piece:

A new network "philosophy and architecture", is replacing the vision of an Intelligent Network. The vision is one in which the public communications network would be engineered for "always-on" use, not intermittence and scarcity. It would be engineered for intelligence at the end user's device, not in the network. And the network would be engineered simply to "Deliver the Bits, Stupid", not for fancy network routing or "smart" number translation.

Fundamentally, it would be a Stupid Network.

In the Stupid Network, the data would tell the network where it needs to go. (In contrast, in a circuit network, the network tells the data where to go.) In a Stupid Network, the data on it would be the boss.

According to Craig Burton, the best geometric expression of the Net's "end-to-end" design is a hollow sphere: a big three-dimensional zero. Across it, every device is zero distance from every other device. Yes, there are real-world latency issues. No path across the void is perfect. But the ideal is clear: the connection between any two computers should be as fast and straightforward as the connection between your keyboard and your screen. Value comes from getting stuff out of the way, not from putting stuff in the way—especially if that stuff is designed to improve performance selectively. The middle is ideally a vacuum. You can improve on it only by making it more of a vacuum, not less. And, like gravity, it should work the same for everybody.

So I see the challenge here as a Search for Terrestrial Stupidity. And I think it's a challenge that goes directly to *Linux Journal* readers and their friends. We are the kinds of people (and perhaps some of the actual people) who imagined and built the open Internet that the whole world is coming to enjoy. And we're the ones who are in the best position to save it from those who want to make it gravy for television.

In other words, we need smart people to save the Stupid Network. I look forward to seeing how we do it. ■

Resources for this article: www.linuxjournal.com/article/9261.

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

THE PENGUIN

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

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

MAY CAUSE DROWSINESS



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

www.penguincomputing.com



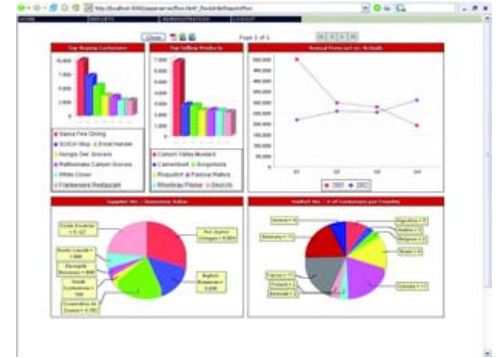
REALLIFELINUX

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

JasperSoft's JasperServer Professional

If you are charged with the task of leveraging IT to help your company make better business decisions, check out the new JasperServer Professional from JasperSoft. Built on the JasperServer Open Source Project, this product is a business intelligence (BI) server that offers ad hoc reporting and analysis intended to simplify the creation of customized reports. JasperServer Professional offers, according to its maker, "everyone in an organization the power to create his or her own BI reports" that are tailored to his or her own needs. In addition, the server is certified with a wide range of third-party platforms, including Apache Tomcat, MySQL, various Linux distros and Unices and more. Customers opting for the subscription service can obtain enterprise-class support and training, indemnification, commercial licensing, access to the customer portal and so on. The Open Source edition of JasperServer, as well as an evaluation edition of the commercial product, are available for download at JasperSoft's Web site.

www.jaspersoft.com and www.jasperforge.org



eXcito's Bubba Server

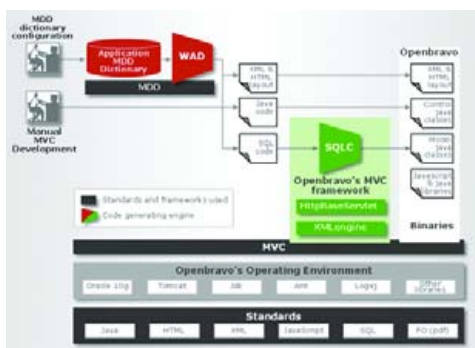
No, friends, this new computer from eXcito has nothing to do with our president emeritus, Bill Clinton! Bubba Server, recently released by eXcito of Sweden, is a diminutive, multifunction, Debian-powered device for the home or SOHO, dubbed by its producer as a "lifestyle home-server". After connecting Bubba to broadband, it's ready to function as any number of servers for you right out of the box: file, Web, FTP, backup, mail (IMAP, SMTP, POP) and so on. Bubba's main features, sayeth eXcito, are the ability to "access your files and different e-mail accounts from any location", its small footprint (18 x 11 x 4cm) and quiet, fanless operation (max. 28dB in active mode). You can acquire your own Bubba with either an 80GB or 250GB hard drive.

www.excito.com

Lenovo's ThinkPad T60p Linux Mobile Workstation

While the heavy-hitting PC makers have shipped and supported Linux desktops worldwide, here at home in the US they have been intimidated into keeping their cupboards bare. What would we do without our scrappy entrepreneurs who have built their Linux-PC empires from scratch? In order for the Linux desktop finally to get traction, the big guys need to bless it and support it, and perhaps Lenovo's new ThinkPad T60p Linux Mobile Workstation will finally ignite some momentum. With the T60p, Lenovo is starting with the high end, as this device is intended mainly for electronic engineers doing integrated circuit and board-level design who desire mobility. Lenovo is currently certifying requisite design apps from companies such as Cadence, Synopsys and Mentor Graphics, which will run on top of SUSE Linux Enterprise Desktop 10 (SLED 10); the latter is fully supported via Lenovo's Help Center. One drawback to the T60p is that, although SLED 10 is supported, it does not come pre-installed. We hope that the efforts invested here will trickle down.

www.lenovo.com



Openbravo r2.11

It is exciting to see a range of firms leveraging the open-source model to provide high-end applications. Openbravo (both the company and app name) finds its niche as an open-source, Web-based enterprise management solution for small and mid-sized enterprises. The application provides for fully integrated management of key business functions such as CRM, billing, data, procurement, inventory, projects, services, production, financial/accounting and business intelligence. Openbravo claims that its architecture is "revolutionary", utilizing "a unique combination of MVC and MDD development frameworks", as well as its own engine for generating application binaries from the MDD dictionary, called WAD. New features in the new r2.11 include several new modules, expanded Web Services features, an improved interface and expanded documentation.

www.openbravo.com



Trolltech's Qtopia Greenphone

Linux companies love colors. We've seen red hats, black ducks, yellow dogs and blue bicuspid. What could be next, pray tell? Trolltech says, green phones! The company's brand-new product, the Qtopia Greenphone, is an open, Linux-based mobile device for application developers of all stripes, allowing them to "create, modify and test Linux-based mobile phone applications on a working GSM/GPRS device" that also has a functioning camera. Trolltech offers Greenphone as part of its software development kit containing the Qt-based Qtopia Phone Edition, an application platform and a UI for Linux-based mobile phones. The company says that Greenphone offers a number of product features and benefits, including an open software stack, accelerated time to market, simplified development processes and reduced costs. Trolltech also sees Greenphone as the first in a series of open mobile devices. Next up, mauve?

www.trolltech.com

Movidis Inc.'s Revolution x16 Server



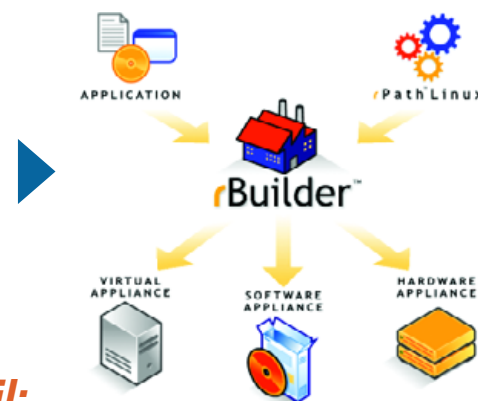
Like Trolltech, Movidis, Inc., has gone "green" as well, only theirs is related to its eco-friendliness. The firm's new Revolution x16 Server is built to provide a single architecture that will perform multiple server functions, for both applications and storage, while consuming a mere 50 Watts. Movidis' approach is to utilize Cavium Networks' OCTEON CN3860, a 16-core, 64-bit MIPS processor that will execute nearly 20-billion instructions per second. According to Movidis, the OCTEON is "optimized for moving data around a network—just what most servers spend their cycles on". The Revolution x16's other key features are integrated accelerators that perform encryption, compression and TCP packet processing in hardware rather than software, as well as Debian burned into the on-board Flash. The Revolution x16 is available in 1U or 2U rackmount enclosures, with either four or eight SATA or SAS drives for a maximum capacity of 6TB on a single 2U RAID system.

www.movidis.com

rPath's rBuilder

The folks at rPath have upgraded their rBuilder product to version 2.0, a platform for creating and maintaining Linux-based software appliances. In essence, rBuilder allows the ISV to import a desired application and combine it with the company's own rPath Linux and create either a software, hardware or virtual appliance image. The result, says rPath, is a reduction in software complexity and cost "by making the operating system disappear". With the appliance as a solution, the ISV can provide its customers with a simplified installation, integration and maintenance process. New in version 2.0 are improved appliance administration, easier customization of administrative interfaces, simplified updates delivered via the Internet to customer locations and the ability to create CD/DVD images for demonstration purposes so that customers can try before they buy.


www.rpath.com



Curtis Smith's *Pro Open Source Mail: Building an Enterprise Mail Solution* and Adrian Holovaty and Jacob Kaplan-Moss' *Pro Django: Web Development Done Right* (both from Apress)

Yes, dear readers, you are getting your money's worth! Because Apress has so many sweet books coming out, your editor failed to pick just one. Book one: Curtis Smith's *Pro Open Source Mail* is a "comprehensive guide to managing the most important mail-related services, including user administration, mail transfer agents, virus protection, spam and mail filtering, Web-based mail and mailing list maintenance". Some applications and tools covered include Sendmail, Qpopper, Dovecot, SpamAssassin, ClamAV and SquirrelMail. Book two: Adrian Holovaty and Jacob Kaplan-Moss' *Pro Django* is a tutorial and reference about the red-hot Django Web development framework. The authors cover everything from creating the components that power a Django-driven Web site to utilizing advanced Django features (such as outputting RSS and PDF and caching). Also included is a range of detailed reference information, such as configuration options and commands. Holovaty is a co-creator of Django; Kaplan-Moss is its lead developer.

www.apress.com



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



Interview with **TIM BRAY**

The loud Atom evangelist Tim Bray talks about everything from Ruby to simplified equal opportunity.

James Gray • Photograph By Johann Wall

No history book on the Internet would be complete without a chapter on Tim Bray. Not only was Tim a co-editor of the XML 1.0 specification, but he also created the first parser software for XML documents and has been co-driving the development of Atom. Today, fulfilling a dual role as tireless Netizen-evangelist and Director of Web Technologies for Sun Microsystems, Tim continues to build on his early work by advocating for a more elegant, platform-independent and user-friendly Internet. *Linux Journal* recently checked in with Tim Bray to get an update on where he is channeling his creative energies these days.

LJ: You have been Director of Web Technologies for Sun Microsystems for just more than two years now. Can you tell us what kinds of projects you've been pursuing in that role?

TB: The most important project is helping return Sun to the position it should be in: profitable and growing. At Sun, I've been a generalist, which is good for someone with adult-ADD. I did a lot of work on launching the employee blogging (see blogs.sun.com). I've been an evangelist in favor of Sun embracing alternatives to Java—both running languages like PHP and Ruby on the Java platform, and embracing those languages in their native form as perfectly viable options for developers. I've been a vocal skeptic of the WS-* project, preferring simpler, more lightweight alternatives based on proven Web technologies. And I've been doing a ton of work on the Atom technology, both co-chairing the IETF working group and evangelizing it to developers. I've also done some work on disk I/O performance (I'm the original author of the venerable "Bonnie" benchmark). Finally, I've been whittling away at a skunkworks named Sigrid for a couple of years now, but <blush> have yet to release anything.

LJ: Does the position at Sun give you an effective "bully pulpit" from which to effect positive change on issues important to you?

TB: No change comes easy, and no single individual has a fulcrum placed in such a way that he or she can move the earth. I've put my weight behind a few ideas and efforts that have moved forward in a way that pleased me, and I've failed to make much progress in some other areas. I think the degree to which I'm listened to has more to do with what I say—whether it makes sense and is interesting—than who I work for.

LJ: What issues and trends are you currently most passionate about, and what form is your advocacy taking?

TB: I think Atom, both the format and the protocol, are going to be pervasive technologies that will have highly visible consequences. Based on my XML experience, I'm now too smart to try to predict exactly what those consequences will be. But I'm evangelizing it everywhere I get a chance, most recently from the stage at OSCON. Enough others have taken up the task of questioning WS-* that I no longer feel compelled to speak up quite so often. Second, I am a very small part of the groundswell of developers heading in the direction of dynamically typed languages. I am personally quite passionately convinced that almost all DRM technologies are technically broken and bad for business, but this

has little to do with my day job.

For all the things I care about, I find my blog (www.tbray.org/ongoing) the most effective way to share my views with the world; mostly because it's a conversation, not a bully pulpit.

LJ: As you say, you have been devoting a great deal of energy to Atom. What is your specific role in it, and what are your thoughts on where it is headed?

TB: I'm the co-chair of the IETF Working Group, and one of the loudest Atom evangelists. Both the Atom data format and the Atom Publishing Protocol (APP) are going to be big. The data

"I AM PERSONALLY QUITE PASSIONATELY CONVINCED THAT ALMOST ALL DRM TECHNOLOGIES ARE TECHNICALLY BROKEN AND BAD FOR BUSINESS, BUT THIS HAS LITTLE TO DO WITH MY DAY JOB."

format will be used in some places where RSS is now, but it turns out that there is a demand for a general-purpose "collection" format—something XML has never had, and Atom suits that bill. The APP provides a low-friction, simple, standardized way to post anything (words, pictures, movies) to the Web, to update it, and to delete it. There's an excellent chance that it will be included in a high proportion of the future's cell phones, not to mention e-mail and Web and news and office-productivity clients, which will thus be able to post to any Web-publishing service that plays by the rules.

LJ: You were one of the three editors of the original XML 1.0 specification. What are your thoughts on the results eight years later?

TB: I'm horribly unsatisfied and keenly aware of all the ways in which XML could have been better, mostly by being smaller and simpler. XML addressed a huge, painful

problem (standardized machine-independent data format) at the right time, and it didn't suck just enough, so it became the default solution. That aside, I'm happy that the world has bought into the notion of sending data around in a way that is thoroughly internationalized and radically independent of any programming language or operating system or hardware.

LJ: From what I understand, XML grew out of the unwieldy SGML and a project to put the Oxford English Dictionary (OED) on-line, is that correct?

TB: Not quite. XML grew out of SGML, which was used in a lot of high-end publishing systems, but not the New OED project, which I managed. The electronic OED, at the time I was there, used a markup system that was an awful lot like what we now call XML; as a side effect of working with it, and then of founding a company, Open Text, to take our inventions to market, I became familiar with SGML. In 1996, when Jon Bosak was getting the XML Project launched, there were maybe a dozen people in the world who were familiar with both SGML and with the Web, and I was one of them.

LJ: How do you think things stand regarding XML as a document format, post-controversy with the state of Massachusetts regarding the Open Document Format?

TB: The fight is basically over; the public sector has noticed the risk reduction and flexibility you get from storing long-lived documents in an XML-based format and has started a move in that direction that will become pervasive, worldwide. After years of trying to convince them that proprietary formats for public data were okay, Microsoft has shifted gears and is frantically trying to apply a thin coat of "standards" paint to its own Office XML formats; the spec is thousands of pages in length and will never be fully implemented by anything but Microsoft Office, which kind of misses the point. I'm pretty confident that the public-sector policy-makers will see through this pathetic ruse. The private sector, which typically has a shorter temporal horizon, is less far along the road to taking good care of its information resources, but it'll get there too.

LJ: Can you tell us more of your thoughts on the state of Web Services?

TB: It depends what you mean by "Web Services". There are two core ideas, taken from the Web (hence the name). First, instead of trying to define APIs across networks, you specify the messages that are exchanged, and second, that XML is a decent data format for the messages. I personally believe both, and this is the only sane way to integrate applications in a heterogeneous environment. Unfortunately, the attempt to standardize this excellent idea, under the WS-* label, has gone off the rails. It is insanely complex, baroque and abstracted, and implementing it requires reading hundreds of pages of poorly written, unstable specs and dealing with ferocious inter-vendor politics.

“SO, I EXPECT WS-* TO FALL FAR SHORT OF EXPECTATIONS, BUT WEB SERVICES, DONE MORE SIMPLY, TO BE THE DEFAULT WAY OF DOING THINGS IN THE FUTURE.”

Fortunately for the future, there is a rock-solid base of proven, efficient, scalable, standardized technology: HTTP, XML and so on, and some very clear guidance on how to do things: REST. So, I expect WS-* to fall far short of expectations, but Web Services, done more simply, to be the default way of doing things in the future.

LJ: What is your take on the Java/OSS debate?

TB: I'm not 100% convinced that an OSS license for Java will bring that much engineering benefit. One of the biggest benefits of open source is that bugs are caught and fixed more quickly. But "Java" is defined as "a binary that passes the TCK", and that's worked well; the Java community likes it, and that's how it's going to stay. So I'm not sure that it's reasonable to expect the open-source "release early, release often" culture to come to Java. On the other hand, Java's licensing has been a cultural obstacle to a lot of people, especially in the Linux space. One result is that things like GNOME and KDE are still substantially written in C++, blech. So I'm optimistic that a real open-source license will eventually empower the developers of the non-Microsoft desktop.

LJ: You are a big fan of Ruby and Ruby on Rails. What is it about them that interests you?

TB: In fact, I'm a fan of dynamic languages in general—my own Weblogging system is written mostly in Perl. To my eye, Ruby and Python stand out from the crowd of such languages in that they seem useful for building large, ambitious software projects as well as the quick one-offs that "scripting" languages have traditionally been used for. Speaking

personally, I find Ruby a bit more pleasing than Python, but the margin is at best 55/45; there are areas where Python is more attractive. I don't think either is going to wipe out the other. Rails seems special; in its sweet spot, maintainable Web apps with a low barrier to entry, it seems like it's set a new standard.

LJ: What is your take on the state of PHP?

TB: Aesthetically speaking, I don't like PHP. I am told that its top-level namespace has 5,000 functions, which is sort of mind-boggling. On the other hand, I've seen how it's empowered legions of people, many without a lot of formal training, to get very usable Web apps on the air quickly. And its scaling story is impressive: anything that runs the infrastructure for Yahoo! Finance deserves respect. Still, whenever I'm asked to look at actual PHP code, chances are it'll be an unmaintainable mess: spaghetti SQL wrapped in spaghetti PHP wrapped in spaghetti HTML. I think that what we'd like, ideally, is something that has PHP's ease-of-use and scaling advantages, but is more effective at separation of concerns and maintainability. Something like Rails or Django. And, Java EE is moving in that direction fast with release 5.

LJ: The Weblog you mentioned above at www.tbray.org/ongoing—what is your mission with it?

TB: No mission whatsoever. I like being able to talk to the world, and even more, I like having the world talk back to me. I'm naturally a fast writer with lots of strong opinions, and it turns out that (for the last couple of years anyhow) a lot of other people are interested in the same things I am. It also gives me a place to post my pictures and talk about politics and music and books and so on. I'll confess that there have been a few occasions when I've deliberately tried to write something to appeal to a big audience or make an impact, and it never works. I totally can't predict which of my pieces will get Slashdotted and which will sink without a trace. So I just write about what I care about. Sometimes people ask me to write about something—sometimes people at Sun, sometimes from elsewhere; sometimes I say yes, sometimes no, based on whether it's interesting or not.

LJ: The potential for everyone to participate in the Internet experience seems to be an important issue for you. In fact, you've written that "The Net itself is a contribution, by humanity to humanity, the engine of future contribution and experience." What do you think it will take to make your vision of a truly accessible and egalitarian Internet into reality?

TB: I'm 100% in favor of making the Net "accessible", but I don't think either the Net or the world are particularly egalitarian. The only equality you can hope for, really, is equality of opportunity. Take blogging for example. Not everyone likes writing, not everyone writes well and not everyone writes quickly. Having said all that, I think that a lot more people could be participating than are right now, and the biggest barrier to entry is the lousy quality of the tools. I think that to improve the quality of the creative experience, we need to get some standard protocols in place, which is why I'm so enthused about the APP.


LJ: All right, since this is *Linux Journal*, we need to ask at least one pure Linux question before we close, okay? What is your favorite Linux distribution?

TB: My favorite distro is Ubuntu, although the server/firewall box in my basement is basic Debian, and I'm happy with that too.

LJ: Thank you for your great insights, Tim! ■

James Gray is Products Editor for *Linux Journal*.

What is the 5-letter word for "an Inexpensive Compact Panel PC"?



The PPC-E5 is the Solution that Fits !

Setting up a Linux Panel PC can be a *Puzzling* experience. However, the PPC-E5 comes ready to run with the Operating System installed on flash disk. Apply power and watch the Linux X Window Graphic User Interface appear on the vivid color LCD. Interact with the PPC-E5 using the responsive integrated touchscreen. Everything works out of the box, allowing you to concentrate on your application rather than building and configuring device drivers. Just Write-It and Run-It. Starting at \$995.

Since 1985
OVER
21
YEARS OF
SINGLE BOARD
SOLUTIONS

EMAC, inc.

Phone 618 - 529 - 4525 Fax 618 - 457 - 0110
2390 EMAC Way, Carbondale, Illinois 62901
World Wide Web: <http://www.emacinc.com>



Their idea of a custom hardware solution is sprinkles on top.

At Pogo Linux we get it-your business and IT needs are unique. So why settle for cookie cutter hardware solutions when you can have exactly what you want and when you need it. Now featuring a complete line of rackmount servers and storage appliances built with Dual-Core Intel® Xeon® Processors.



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

To get started, contact us at 888.828.POGO or inquiries01@pogolinux.com
Pogo Linux, Inc. 701 Fifth Ave. Suite 6850, Seattle, WA 98104

www.pogolinux.com





Caller ID with. Asterisk and Ajax

COMBINE ASTERISK AND AJAX TO
DISPLAY INCOMING AND OUTGOING
CALL INFORMATION. • MIKE DIEHL

I've been using an Asterisk server to handle all of our telephone service for about a year now. During this time, I've discovered many really neat things that can be done with Asterisk, VoIP and various other technologies. One of the more gimmicky things I've done is sent the caller-ID information from incoming calls to a Web page on my browser, in real time. To do this, I had to use Asterisk, Perl, CGI, HTML, CSS, SQL, XML and Asynchronous JavaScript, or Ajax. There are a lot of different pieces to bring together, but sometimes that's what makes a project interesting.

Here's how it works in a nutshell. When someone calls us at the house, the Asterisk server waits for the caller-ID information to be sent. The server then puts this information, and a few other pieces of information, into a file in a subdirectory under /tmp. This is all done in the Asterisk dial plan. Then, I have a Web page open in my browser that runs a JavaScript program every second. This JavaScript program uses an XMLHttpRequest object to query the server for new caller-ID information. The CGI script on the server returns an XML file containing the caller information. The JavaScript program parses the returned XML and displays the content. I've created a Cascading Style Sheet (CSS) that makes the caller information look like a sticky note placed on the Web page. When the incoming call is complete, the Asterisk server creates a Call Detail Record, or CDR, which resides in an SQL database.

Each time the JavaScript contacts the server, the CGI script looks for the CDR. If it exists, the program knows that the call is over and deletes the caller information file in /tmp. This has the effect of causing the sticky notes to disappear when the call is complete.

As an added bonus, the program supports up to four concurrent calls and can be used to indicate outbound calls as well. It's kind of nice to be able to see who's on the phone, regardless of whether the person is the caller or callee, without having to interrupt the person on the phone to ask. When my boys get older, this may become an even more important feature.

For this system to work, you must configure your Asterisk server to put CDRs in an SQL database. By default, Asterisk puts CDRs in a comma-delimited file. The problem is that the flat file CDRs don't contain the call's unique ID, which this system uses to detect when a call has completed. The CDRs that get put into the SQL database contain this field. This shouldn't be a steep requirement though. As I recall, configuring Asterisk to store CDRs in a Postgres database was fairly

straightforward and well documented in the cdr_pgsqldb.conf file. You also could use a MySQL or ODBC database, if you like.

The first, and easiest, part of this project is to modify the Asterisk dial plan to create the flat file when an incoming or outgoing call is made. Once you determine where to make the change, it's a simple

Listing 1.

Example Web Page

```
<html>
<head>
  <title>CID Test</title>
  <script language=javascript src=http://hostname/cid.js>
  </script>
  <style type="text/css">
    @import "cid.css";
  </style>
</head>
<body>
  <div id="phone1"></div>
  <div id="phone2"></div>
  <div id="phone3"></div>
  <div id="phone4"></div>
  <script>
    start_cid();
  </script>
  Your Content Would Go Here.
</body>
</html>
```

Listing 2.

Sample cid.css File

```
div#phone1{
  background: #FFFFCC;
  display: none;
  position: absolute;
  border-top: thin solid black;
  border-left: thin solid black;
  border-right: 6px solid black;
  border-bottom: 6px solid black;
  top: 85%;
  left: 2%;
  width: 20%;
  height: 5em;
}

div#phone2{
  background: #FFFFCC;
  display: none;
  position: absolute;
  border-top: thin solid black;
  border-left: thin solid black;
  border-right: 6px solid black;
  border-bottom: 6px solid black;
  top: 85%;
  left: 27%;
  width: 20%;
  height: 5em;
}

div#phone3{
  background: #FFFFCC;
  display: none;
  position: absolute;
  border-top: thin solid black;
  border-left: thin solid black;
  border-right: 6px solid black;
  border-bottom: 6px solid black;
  top: 85%;
  left: 52%;
  width: 20%;
  height: 5em;
}

div#phone4{
  background: #FFFFCC;
  display: none;
  position: absolute;
  border-top: thin solid black;
  border-left: thin solid black;
  border-right: 6px solid black;
  border-bottom: 6px solid black;
  top: 85%;
  left: 77%;
  width: 20%;
  height: 5em;
}
```

AS AN ADDED BONUS, THE PROGRAM SUPPORTS UP TO FOUR CONCURRENT CALLS AND CAN BE USED TO INDICATE OUTBOUND CALLS AS WELL.

one-line addition, as shown here (all one line):

```
exten => s, n, system(echo "IN${CALLERID(name)}
➡#${CALLERID(number)}#${UNIQUEID}" >
➡/tmp/panels/cid/${UNIQUEID})
```

This line creates a file in `/tmp/panels/cid` that contains four fields, delimited by the `#` character. Of course, you need to create `/tmp/panels/cid` and give it appropriate permissions so that the Asterisk server can create files in it and the CGI script can read and delete those files. The first field is either `IN` or `OUT` and indicates that the call is `INcoming`, or `OUTgoing`. The next two fields call the `CALLERID()` function to retrieve the caller's name and phone number. The last field is the call's unique identifier. You need to place this line in your dial plan, such that the server has already received the caller-ID information but before the call is handed off to the dial command. If you want to receive information about outgoing calls, you could add a

line like this to your dial plan:

```
exten => s, n, system(echo "OUT##${EXTEN}#${UNIQUEID}"
➡> /tmp/panels/cid/${UNIQUEID})
```

In the case of the outgoing call, we don't have any caller-ID information to display, so the second field is left blank. We do know the number that was dialed, which is retrieved via the `${EXTEN}` variable in the third field.

In both the incoming and outgoing cases, you need to make sure to update the extension field and the priority fields (`s` and `n` in this example).

For the purpose of demonstration, I've stripped the Web page down to its most basic requirements, as shown in Listing 1.

This seemingly simple HTML code does a lot of things. First, it loads the `cid.js` JavaScript code. Then, it imports a stylesheet called `cid.css`. This stylesheet will give you a lot of flexibility to customize the appearance of the sticky notes. Then, the HTML code creates four div sections, called `phone1` through `phone4`. These sections will be made

Listing 3.

CGI Script

```
#!/usr/bin/perl

use DBI;

$dbh = DBI->connect("dbi:Pg:dbname=database", "postgres", "password")
|| die "Can't connect to database.\n";

print "Content-type: text/xml\n\n\n";

print "<panels>\n";

check_cid("/tmp/panels/cid");

print "</panels>\n";

exit;

sub check_cid {
    my($dir) = @_;
    my(@a, $a, $file, $count, $stop);
    local(*FILE, *DIR);

    opendir DIR, "/tmp/panels/cid";
    while ($file = readdir(DIR)) {

        if ($file eq ".") { next; }
        if ($file eq "..") { next; }

        open FILE, "/tmp/panels/cid/$file";
        chomp($line = <FILE>);
        close FILE;

        ($dir, $name, $number, $uid) = split("#", $line);

        $count++;

        if ($dir eq "IN") {
            $html = "Incoming call from $name ($number)";
        } else {
            $html = "Outgoing call from $name ($number)";
        }

        expire_call($uid);

        print <<EOF
<panel>
  <name>phone$count</name>
  <content>$html</content>
</panel>
EOF
;
    }
}

sub expire_call {
    my($id) = @_;
    my($sth, $count);

    $sth = $dbh->prepare("select count(*) from cdr where
uniqueid=\'$id\'");
    $sth->execute();

    ($count) = $sth->fetchrow_array();

    if ($count) {
        unlink("/tmp/panels/cid/$id");
    }
}
```

Listing 4.

Resulting XML File

```
<panels>
<panel>
  <name>phone1</name>
  <content>Incoming call from Mike Diehl (15055558592)</content>
</panel>
</panels>
```

visible later on and will be filled in with caller information. Finally, the HTML code starts the periodic polling by calling the `start_cid()` function. We'll discuss that function later.

Even though my CSS skills aren't world-class, I've included a sample `cid.css` file to get you started (Listing 2).

This CSS file could have been made more concise by putting all of the common formatting in a common class; I'll leave that as an exercise for the reader. This stylesheet creates four evenly spaced sticky notes at the bottom of the screen. The sticky notes are yellow with a neat 3-D drop-shadow effect (Figure 1).

Now, it's time to take a look at the CGI script (Listing 3).

This Perl script scans the `/tmp/panels/cid` directory for files, skipping the `.` and `..` entries. Each file it finds is opened and read. The final result is an XML file like the one shown in Listing 4.

Of course, the XML file could contain up to four `<panel>` blocks corresponding to `phone1` through `phone4`. The `<content>` block contains the text that is put into each sticky note. I've found that because this is

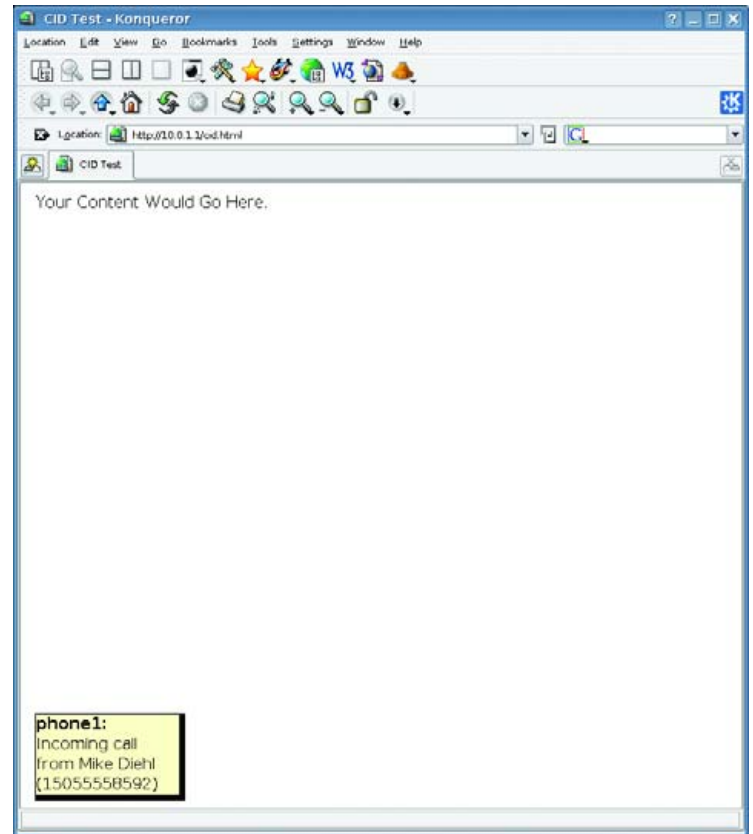


Figure 1. The incoming call information is displayed in a Web page in sticky note format.



Expert Included.

Matt provides expert, dedicated technical support, backed by years of training and experience. One reason he's so good at his job is that he feels a personal connection to customers: when they're happy, he's happy. Matt likes systems based on the Next-Generation AMD Opteron™ processor because he knows they offer a seamless upgrade path to Quad-Core while maintaining the same thermal envelope. Also, AMD Virtualization™ allows testing of new operating systems and applications, with painless rollback until everything is just right, and that makes Matt happy. After all, Matt would rather read about new technology than repair it.

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



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

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

an XML file, it's difficult to embed HTML in the <content> block, so I don't do much formatting of this text. It's fairly easy to see how incoming and outgoing calls are handled separately.

As the XML is generated for each phone call and sent to the client, the call to `expire_call()` is made. This function simply searches the CDR database to see if the phone call has been completed. Asterisk adds CDR records only when a call is concluded, so if the record is in the database, the call is finished and the file in `/tmp/panels/cid` can be removed.

The JavaScript component is both the workhorse of the system and the most difficult part to understand (Listing 5).

As mentioned previously, the whole system is started by the initial call to `start_cid()`. All this function does is arrange for the `update_cid()` function to be called every second. The `update_cid()` function makes a call to `get_from_server()` to get an XMLHttpRequest object in a browser-independent fashion. This request object is returned for later use.

Next, the `update_cid` function calls `clear_panels()`, which simply arranges for each sticky note to be empty and invisible, initially. The sticky notes will become visible as we put content into them.

The rest of the program is a bit more difficult to follow. Using the request object mentioned earlier, and the `getElementsByTagName()` function, we get an XML object with the <panels> block intact. Another application of the `getElementsByTagName()` applied to this XML object gives us an array of individual <panel> blocks.

Then, we start a loop over each <panel> block in the array with the understanding that each time through the loop will correspond to a phone call in progress; we'll create a new sticky note for each call. Each <panel> block contains a <name> and a <content> block, the values of which we extract into appropriate variables. Then, by using the `getElementById()` document method, we find the <div> element in the HTML document with the same ID as the name of the panel. Now we have all of the information we need about the sticky note: the name, the content and the location in the Web page. So, we set the <div> block to be visible, then assign some content to it via the `innerHTML` attribute. Finally, we go back to the top of the loop and continue again.

This "poll the server and display the results" process runs every second without any intervention from the user and without having to reload the Web page. This gives the user the perception that the sticky notes simply pop up when the phone rings and disappear when the phone is hung up.

As you can see, JavaScript is a very powerful language. Unfortunately, browser support and development tools for JavaScript are poor to nonexistent. During the development of this program, I had to contend with browser crashes, inadvertently cached information and cryptic runtime error messages. Once I got it working, I had to make sure it worked on each of the browsers I use regularly, Konqueror and Firefox. I suspect that it will run on "that other browser", but I've not tested it. Because I do most of my software development with vi, I'm not really big on Integrated Development Environments (IDEs), but if you know of one that works well for JavaScript, I'd love to hear from you.

Now that the program is working, it's time to think about ways to improve and extend it. The first obvious change I'd like to make to this program is to have it display a hyperlink that would allow me to bring up additional information about the caller. It could get this information from my contact list or even from an additional database. Maybe it could display a picture of the caller, though it might take a lot of time to photograph all my friends, family and acquaintances. It might also be nice to have a button display for incoming calls that would allow me to reject an incoming call and have it go straight to voice mail. I could also extend this same method to have a Web page display other information besides caller ID. It wouldn't be hard to extend this system to let me know when I have unread voice mail waiting, or when my friends become available for chat via IM.

So there you have it—a fun little toy that brings together many different tools and technologies. Recalling that Qwest used to charge us \$6 US a month for caller ID, I wonder what they would charge to make it Web-accessible? ■

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

Listing 5.

The JavaScript Component

```
function start_cid () {
    setInterval("update_cid()", 1000);
}

function update_cid () {
    var req;
    var xml;
    var panels;
    var count;
    var name;
    var div;

    req = get_from_server();

    clear_panels();

    xml = req.responseXML.getElementsByTagName("panels")[0];

    panels = xml.getElementsByTagName("panel");

    for (count=0 ; count < panels.length ; count++) {
        panel = panels[count];

        name = panel.getElementsByTagName("name")[0];
        name = name.firstChild.nodeValue;

        content = panel.getElementsByTagName("content")[0];
        content = content.firstChild.nodeValue;

        div = document.getElementById(name);

        div.style.display="block";
        div.innerHTML = "<b>" + name + " : </b>" + content;

        if (div.innerHTML == "") {
            div.style.display="none";
        }
    }
}

function get_from_server () {
    var req;

    if (window.XMLHttpRequest) {
        req = new XMLHttpRequest();
    } else if (window.ActiveXObject) {
        req = new ActiveXObject("Microsoft.XMLHTTP");
    }

    req.open("GET", "/cgi-bin/cid.pl", false);
    req.send(null);

    return req;
}

function clear_panels () {
    for (count=1 ; count < 5 ; count++) {
        document.getElementById("phone" + count).innerHTML = "";
        document.getElementById("phone" + count).style.display="none";
    }

    return;
}
```

Polywell Server and PC Solutions

Specialize for Small to Medium Size Business

High-End Servers and Workstations

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

Low-Cost Linux PCs

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

- 2 x AMD® Optreron™ Dual-Core Processors 2218+ with Hyper Transport Technology
- Upto 64GB 667/533MHz ECC DDR2 (16 Sockets)
- Upto 2TB 4 x 500G Swap Drive per 1U Rack
- 1 x PCI-X 133MHz or 1 x PCI-E RISER Slot
- Dual Gigabit Ethernet, ATI Graphics, 4 SATA-RAID
- CD-ROM Drive, Optional DVD-RW or CDRW
- SATA RAID-5 or SCSI Controller on-board
- 1U 24" Depth Rack Chassis with upto 600W PS
- 4 x Swappable Drive Bays (SATA or SCSI)
- Supports Linux, FreeBSD or Windows
- Custom Configuration Available
- Please call for other Options



1114SAS-2500A, 2218+, 32GB, 2TB **\$8,599**
 1114SAS-2500A, 2212+, 16GB, 1TB **\$4,399**
 1112ES-2500A, 2210E, 2GB, 500GB **\$1,999**



4024AIS 4U 12TB SAS/SATA **\$10,500**
 4024AIS 4U 6TB StorageServer **\$6,399**

Linux Appliance PCs

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

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



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

64GB DDR 4-way Workstation

- 2 x AMD® Optreron™ Dual-Core Processors 285+ Upto 64GB ECC DDR (16 Memory Sockets)
- 2 x 133MHz PCI-X, 1 x PCI-E x16 Slots
- Dual Gigabit Ethernet, 4 x SATA-RAID Controller
- Quiet and Cool 11-Bay Tower + 1400W 80Plus P/S
- 2x320G HD, DVD-RW, Floppy, Optional Card Reader
- 8-Layer Motherboard with Special Quiet Cooling
- Optional Sound Card, 1394 firewire
- Supports 64/32-bit Linux, FreeBSD or Windows
- Special for Large Memory Intensive Applications
- Built-to-Order or Configure-to-Order

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



64G RAM, 2x285+, QuadroFX5500 **\$20,999**
 32G RAM, 2x265+, QuadroFX3450 **\$8,499**
 16G RAM, 2x246+, QuadroFX1500 **\$4,650**
DDR2 Socket-F is now available

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

1U Power Saving ISP Server

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



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

Pick the style for your application



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

- 8 or 4 AMD® Optreron™ Dual-Core Processors with Hyper Transport Technology
- Upto 128GB DDR Memory for 16-way (32 sockets)
- Upto 64GB DDR2 Memory for 8-way (16 sockets)
- 4x Gigabit LAN, 8x SATA RAID-5 for 16-way
- 2x Gigabit LAN, 4x SATA, U320 SCSI for 8-way
- 4 x 133/100/66MHz PCI-X Slots for 16-way
- 2 x 133, 2 x 66MHz PCI-X, 1x PCI Slots for 8-way
- On-board ATI Graphics, USB 2.0
- 5U 26" Rack 1300W 3+1 Redundant P/S 16-way
- 2U 27" Rack 700W PFC P/S for 8-way
- Supports Linux, FreeBSD or Windows
- Custom Configuration Available
- Please call for other Options



5U 16-way 865, 32GB, 2TB, 8801T5U **\$19,999**
 2U 8-way 8212, 8GB, 900GB, 8450A **\$7,499**

Mainstream Linux Systems

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

939NV3000TD12LX starts at **\$399**

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



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

888.765.9686

www.Polywell.com/us/LJ



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

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

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

Migrating to Drupal

**Why and how
the Planetizen
Web site
migrated to
the Drupal
infrastructure
for communities.**

Drupal is often mentioned in discussions about blogging tools or Web-based forum software. Sure, you can run a blog or an on-line forum using Drupal, but that is only part of what Drupal can do. Drupal is better described as a framework that provides an infrastructure for on-line collaboration and communities. It can be used to run corporate Web sites, intranets, news portals and many other types of Web sites.

The Drupal Project has its roots in an internal message board system built by University of Antwerp student Dries Buytaert for his student dorm. In 2001, Dries released the software as an open-source project named Drupal (pronounced "doo-puhl"). Others started using Drupal and began contributing to the project. Drupal is built using open-source technologies: the PHP programming language and the MySQL or PostgreSQL databases. Licensed under the GNU General Public License (GPL), Drupal can be downloaded and used for free. As with many successful open-source projects, Drupal is maintained and developed by a thriving user and development community. Five years old in January 2006, Drupal has evolved into a robust content management platform.

Working at a Web development firm, we have successfully built many Web sites for our clients based on Drupal. In this article, we share what we have learned, and we tell the story of our most complex Drupal project to date. ▶▶

ABHIJEET CHAVAN AND MICHAEL JELKS



A Drupal Migration Story

Planetizen is a community Web site for urban planners, architects, developers, environmentalists and other professionals. It offers daily news summaries, editorials, jobs and many other services. Launched in 2000, Planetizen has grown into a popular Web site with a large international audience. To manage a constantly updated Web site, such as Planetizen, a content management system (CMS) is a must. We had built our own custom CMS using PHP and MySQL in 2000. As the Web evolved, we wanted to add new features, but doing so meant expensive in-house development. So, we began looking at alternatives.

By this time, numerous open-source CMS projects had matured and offered many of the features we wanted to add. Migrating to a pre-built open-source CMS made sense. We could cut down on development time, add the features we needed and benefit from all the advantages that come with using open-source software. Because we already had experience using PHP and MySQL, we searched for open-source CMSes built using those technologies. After evaluating and testing several different packages, we selected Drupal. (See “Seven Criteria for Selecting Open Source Content Management Systems” in the on-line Resources.)

Why We Selected Drupal

Drupal has many of the features you would expect from a modern CMS, such as user management; access control; work flow; separation of content, presentation and logic; and Web-based editing and administration. Drupal appealed to us for many reasons—here are the top five:

- 5) Sensible URLs and URL aliasing: many CMSes generate long, convoluted URLs that are difficult to share via e-mail or over the phone. Drupal arguably generates the sleekest URLs in the CMS world. Most Drupal URLs are in the format `http://www.planetizen.com/node/156`. Also, Drupal’s URL aliasing feature makes it is easy to create URLs that make sense to readers. Using URL aliasing, the above URL can be mapped to `http://www.planetizen/about/faq`.
- 4) Syndication and aggregation: community Web sites, such as Planetizen, benefit from information flowing in and out of the site. Content stored in Drupal easily can be syndicated to readers or other Web sites using RSS feeds. Also, a news “aggregator” to pull in syndicated content via RSS feeds is built in to Drupal.
- 3) Modular architecture: Drupal’s functionality is organized into modules that can be switched on and off. This approach makes it possible to build different kinds of Web sites with Drupal. If we were going to invest a lot

of time into learning a CMS, it might as well be one that can be adapted for other projects as well.

- 2) Developer-friendly: we anticipated the need to customize any CMS we selected. We felt comfortable with Drupal’s elegantly designed architecture and the consistency of the code. It was relatively easy to understand a feature and start making modifications. Features such as the devel module that displays database queries and variables for each page later proved to be invaluable in migrating to Drupal.
- 1) Taxonomy: our single-most important reason for selecting Drupal was its powerful taxonomy system for categorizing content. It is possible to create a set of descriptive terms and associate content with those terms. The taxonomy system makes it possible to adapt Drupal for a diverse set of content management needs.

Drupal Basics

You can download the latest stable release package from the Drupal Web site. Installing Drupal is a fairly straightforward process. It involves creating a MySQL database, importing tables, copying files, setting file permissions and editing a configuration file. Most of the Drupal options can be configured using its Web-based administration interface. Refer to the `INSTALL.txt` file available with the downloaded package for detailed installation instructions. Additional configuration instructions are available on the Drupal Web site.

In Drupal, most of the content is stored as a node. A node could be a page, a poll or one of the many node types. For example, the page node has a title, body, author, date and some basic attributes. Some modules provide their own node types, which may have additional attributes.

The visual presentation of content is controlled by a theme. Drupal comes with a selection of themes, and it is easy to create your own. Most themes have a central content column and left and/or right sidebar columns. Sidebars can contain blocks of information. Filters control the input format used to store

Name	Description	Enabled
aggregator	Aggregates syndicated content (RSS and RDF feeds).	<input type="checkbox"/>
archive	Displays a calendar for navigating older content.	<input type="checkbox"/>
block	Controls the boxes that are displayed around the main content.	required
blog	Enables keeping an easily and regularly updated web page or a blog.	<input type="checkbox"/>
blogapi	Allows users to post content using applications that support XML-RPC blog APIs.	<input type="checkbox"/>
book	Allows users to collaboratively author a book.	<input type="checkbox"/>
comment	Allows users to comment on and discuss published content.	<input checked="" type="checkbox"/>
contact	Enables the use of both personal and site-wide contact forms.	<input type="checkbox"/>
drupal	Lets you register your site with a central server and improve ranking of Drupal projects by posting information on your installed modules and themes; also enables users to log in using a Drupal ID.	<input type="checkbox"/>
filter	Handles the filtering of content in preparation for display.	required
forum	Enables threaded discussions about general topics.	<input type="checkbox"/>
help	Manages the display of online help.	<input checked="" type="checkbox"/>

Figure 1. Activating Drupal Modules

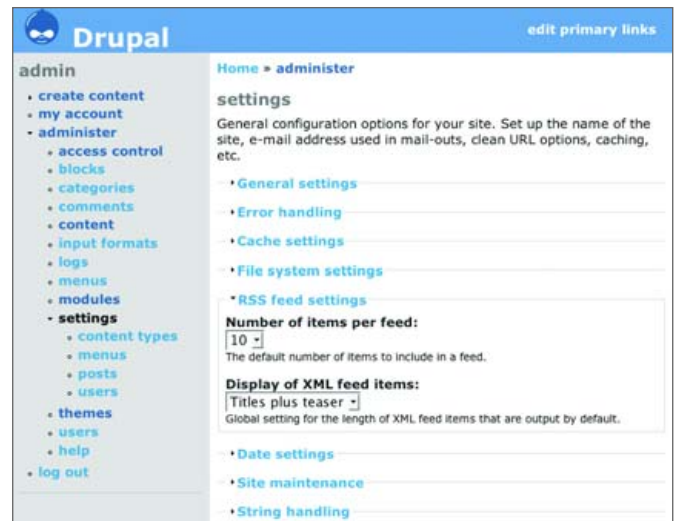


Figure 2. Configuring Drupal

text in nodes or blocks. For example, you can store content in filtered HTML, which limits the HTML tags that can be used. You even can store PHP code snippets.

Five Key “Ingredients”

The basic Drupal install leaves you with a usable Web site to which you can start adding content immediately. But, what you see after installation is only the core functionality. Drupal offers much more. In most cases, you will want to tailor Drupal to your particular content management needs. This is where Drupal’s flexibility can become overwhelming. After building several Web sites with Drupal, we believe the key to creating successful Drupal implementations—“recipes” if you will—lies in understanding the interplay of five Drupal “ingredients”: module selection, configuration, access control, taxonomy and theme.

Module Selection

A module is additional code that extends Drupal's functionality. Drupal comes with a set of core modules, and additional modules can be downloaded and installed as needed. The Drupal Web site lists a large collection of contributed modules created by the community. If you need a particular feature, look for a module that offers it. Several modules may offer similar features or even different implementations of a single feature (Figure 1).

Configuration

By changing configuration options for individual modules and site settings, you can substantially alter the way Drupal behaves. Many modules add features in blocks that appear in a node's sidebar. Often a particular CMS behavior or work flow that you need may just be a matter of configuring modules in a certain way. Be prepared to spend some time experimenting with different settings (Figure 2).

Access Control: Roles and Permissions

Accounts allow you to control what users can see and do on a Drupal Web site. The first user account is considered to be a root account with complete administration privileges. For the other users, you can set what they can do by assigning them to roles. Drupal comes with two roles: anonymous user and authenticated user. You may want to add additional roles, such as editor or manager, and specify what those roles can do. A user can be associated with one or more roles (Figure 3).



The screenshot shows the 'access control' page in Drupal. It has tabs for 'permissions', 'roles', and 'access roles'. Below the tabs is a paragraph explaining permissions. The main part of the page is a table with columns for 'Permission', 'anonymous user', 'authenticated user', and 'editor'. The table lists permissions for various modules like 'block module', 'comment module', 'filter module', 'menu module', and 'node module'.

Permission	anonymous user	authenticated user	editor
block module			
administer blocks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
use PHP for block visibility	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
comment module			
access comments	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
administer comments	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
post comments	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
post comments without approval	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
filter module			
administer filters	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
menu module			
administer menu	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
node module			
access content	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Figure 3. Setting Permissions for User Roles

Taxonomy

Drupal's taxonomy system enables you to associate a node with one or many descriptive terms. You can create multiple sets of terms called Vocabularies. Vocabularies can be flat or hierarchical lists. For each vocabulary, you can specify which node type it applies to. This combination can help you create a classification system for content that suits your particular information architecture

www.faircom.com/go/?track

Your packages get

tracked & delivered

on time by FedEx.

FairCom database technology makes it possible.



FairCom[®]

needs. Many other features and modules depend on the taxonomy. For example, you can generate navigation elements, control access to content or switch visual presentation based on taxonomy. Take the time to develop good taxonomy vocabularies and design them so you can expand them easily in the future (Figure 4).

Theme

Drupal allows you to customize the layouts of pages easily using an extensible theme system. A convenient way to build a custom theme for your Web site is to base it on one of the themes packaged with Drupal. You can use different themes for certain users or in association with taxonomy terms.

Various combinations of the above five ingredients will result in surprisingly diverse solutions. Search the Drupal Web site for “recipes”. If you still cannot achieve what you need, you can customize Drupal or build custom modules.

Migrating Planetizen

We started Planetizen’s migration by making a list of all the features we would need and identifying which Drupal modules would provide that functionality. This required testing different modules and configuration settings. We identified requirements that could not be met using Drupal modules. These features would require custom development. We then developed the taxonomy, defined user roles and permissions, and decided on the work flow. To maintain the original look and feel in the Drupal-based version, we developed a custom theme. Moving to a new CMS is also a good time to rethink current business logic and improve it. We took this opportunity to prune out less-popular Web site features.

The biggest migration challenge was pulling in five years’ worth of data into Drupal—more than 15,000 news stories. Drupal story and page node types provided only basic title and body attributes for a node. Each news item stored in Planetizen had several other attributes. What we needed was our own custom content type. Drupal’s flexinode provides an easy way to create custom content types without programming. Unfortunately, it turned out that the flexinode route would be an inefficient solution for us. Using flexinode, each Planetizen news story would have taken up to eight separate table inserts as opposed to the standard single insert, due to the way flexinode stored data.

Drupal’s wealth of third-party modules came to the rescue. We discovered that a book review module was very similar to what we needed. By examining its code, we were able to customize the book review module to create the content types we needed. We then created custom scripts to insert Planetizen’s data into the appropriate fields directly in Drupal’s MySQL tables.

Limitations and Workarounds

We did encounter some limitations with Drupal. One limitation was the mechanism for maintaining time zones and daylight savings time in Drupal. Our workaround was to use only the PST/PDT time zone and manually update the time zone when it was time for a daylight savings time change. This is a known issue and is being addressed by developers.

Flexinode makes it possible to create custom node types without programming, but as we discovered, it has its limitations. The alternative is to develop custom node types as modules. Drupal provides a solid foundation for creating your own modules, but it requires programming experience. The Drupal team is addressing this issue with the Content Construction Kit (CCK), an effort currently under development that aims to make it easier to create custom node types.

The screenshot shows a taxonomy management interface. At the top, there are tabs for 'list' and 'add vocabulary'. Below this, a text block explains the taxonomy module's functionality. A table lists various categories and subcategories, each with an 'edit term' link. The categories include World, Africa, Asia-Pacific, Australia, China, South Asia, India, Central and South America, Europe, United Kingdom, Middle East, North America, United States, Alabama, Alaska, and American Samoa.

Name	Type	Operations
World	article, news	edit vocabulary add term preview form
-- Africa		(edit term)
-- Asia-Pacific		(edit term)
---- Australia		(edit term)
---- China		(edit term)
-- South Asia		(edit term)
---- India		(edit term)
-- Central and South America		(edit term)
-- Europe		(edit term)
---- United Kingdom		(edit term)
-- Middle East		(edit term)
-- North America		(edit term)
---- United States		(edit term)
----- Alabama		(edit term)
----- Alaska		(edit term)
----- American Samoa		(edit term)

Figure 4. Planetizen’s Taxonomy

The screenshot shows the Planetizen website before migration. The layout is cluttered with multiple columns of content. The top navigation bar includes links for Home, About, Newsletter, Advertise, Search, News, Announce, Op-Ed, TechTalk, Jobs, Consultants, Radar, Sites, Books, and Courses. The main content area is divided into several sections: Consultants, Jobs, Featured Jobs, Announcements, and Featured Announcements. Each section contains various articles and links, such as 'Katrina: An Unnatural Disaster' and 'Learning From Grand Forks'.

Figure 5. Planetizen before Migration

The screenshot shows the Planetizen website after migration to Drupal. The layout is cleaner and more organized. The top navigation bar is simplified, including links for Home, News, Op-Ed, TechTalk, Announce, Jobs, Consultants, Courses, and a search box. The main content area is divided into several sections: User Login, Donate, Consultants, Announcements, and Featured Announcements. Each section contains various articles and links, such as 'Katrina: An Unnatural Disaster' and 'Learning From Grand Forks'.

Figure 6. New Drupal-Based Planetizen

One problem we ran into had nothing to do with Drupal. Our production Web server was running an older version of PHP that could not be upgraded, due to some hosting restrictions. This caused the search module to fail; however, we were able to circumvent this problem by modifying the search module. We thanked ourselves once again that we were using an open-source CMS.

Security patches and core code updates for Drupal are released on a regular basis. This is a good thing, but upgrading customized Drupal installations can be cumbersome. We recommend limiting customizations to specific modules or developing custom modules. Also, using a version control system, such as CVS or Subversion, can help in tracking your customizations against official Drupal releases.

We launched the new Drupal-based Planetizen Web site in September 2005 and received positive feedback from readers. Since the launch, we were able to add new sections and features without having to develop them from scratch (Figures 5 and 6).

Drupal's Future

As we write this article, Drupal's next release, version 4.7.0 is in beta. Improvements include a better default theme engine, refined search functions, improved PostgreSQL support, themeable forms, Ajax-enhanced administration interface and a better upgrade script. Also promising is the development of the CCK that could, along with actions, workflow and views modules, make Drupal even more flexible and powerful.

Some people in the Drupal community predict that the trend to watch in 2006 is the emergence of application-specific Drupal distributions—re-packaged versions of Drupal catering to a particular need. One such distribution is CivicSpace, a community organizing platform popular with grassroots organizations, nonprofits and political campaign Web sites. CivicSpace provides a Web-based installer and a configuration wizard that sets up Web sites for common-use scenarios. It includes a selection of Drupal modules relevant to running community organizing Web sites so you don't have to research, download and install individual modules. CivicSpace also includes CiviCRM, a Web-based constituent relationship management application that offers features, such as on-line fund raising, contact management, tracking volunteers, donors and clients. Efforts are underway to develop similar distributions for educators and artists.

Conclusion

We have used Drupal for several different types of projects, including corporate, collaborative, intranet and academic Web sites. What makes Drupal so versatile?

According to its founder Dries Buytaert, Drupal aims to provide "a solid base to extend and implement custom content management solutions". This may be one of the reasons for its popularity. It strives to be a content management platform that enables developers and users to customize their own unique solutions based on Drupal's core engine. Drupal's modular architecture has resulted in several interesting community-contributed modules. These modules often connect Drupal to other popular programs or services, opening up interesting and unexpected possibilities.

It's true that non-programmers can achieve a lot with Drupal simply by tweaking configurable options. Those with modest HTML or PHP experience can customize themes and layouts or use snippets of code shared by the community on Drupal's Web site. And, of course, PHP experts can create their own custom modules and tweak Drupal as much as they like.

However, its extensibility and flexibility also have made Drupal more complex. The solution you are looking for may be found in a particular combination of modules, configured in a certain way, using a well-crafted taxonomy and carefully thought-out user permissions. Drupal is capable of addressing complex content management needs, but tapping its potential does require a deep understanding of how it works.

What is admirable about Drupal is that it makes it possible—to a certain degree, without writing any code—to shape a diverse range of Web-based solutions built on the same core content management platform. And, it achieves this while remaining true to its stated principles of standards-compliance and collaborative open-source development. Drupal may not have a perfect solution for each problem, but it can meet a lot of different content management needs reasonably well. Ultimately, what matters is that Drupal helps people, whether they are programmers or non-programmers, large organizations or individuals, tap into the collaborative potential of the Web. ■

Resources for this article:
www.linuxjournal.com/article/9264.

Abhijeet Chavan is the Chief Technology Officer of Urban Insight, Inc., a Web development consulting firm. He also is the co-founder and co-editor of Planetizen.

Michael Jelks is a Senior Developer at Urban Insight, Inc., with more than 37 dog years of experience implementing Web-based applications with Perl, PHP and MySQL technologies.

NEW!


TS-7300 High Security Linux FPGA Computer




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

 **200 MHz CPU**

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

 **Design your solution with one of our engineers**

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

 **See our website for options, peripherals and x86 SBCs**



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

SIMPLE WEB SITES USING DOCBOOK XML AND CSS

How to build simple content Web sites using
DocBook XML and CSS. **David Lynch**

The Web was originally intended to make content easily accessible. Today, Web developers focus on style and marketing, but the need to put together content-driven Web sites quickly and easily remains as valid as when Tim Berners-Lee first conceived of HTML. I have taken the approach of using primarily DocBook XML and CSS, as well as some other readily available Linux tools, that allows me to bring up simple content-focused Web sites—a poor man's content management system.

I am an embedded software developer. HTML, XML, CSS and the Web in general are peripheral to what I do. I am not as intimate with the details and idiosyncrasies of HTML as I am of processors, NICs and UARTs. Yet

today, the Web is part of everything. Proof that an embedded processor is up and running under Linux often consists of being able to browse Web pages on it. I look for clients, and clients seek me out over the Web. Although expertise in JavaScript, cross-browser HTML, CSS, PHP, Ruby on Rails and so forth is not essential, a basic knowledge of HTML and the ability to use some tools to create simple but useful Web sites quickly and easily is increasingly a core skill to software development, as well as many other jobs. DocBook XML provides a means of creating documentation focused on content, with the ability to use it easily in many forms, including Web pages.

This approach has a number of elements, and they are not heavily inter-dependent. Even if you do not like my overall approach, you can take bits and pieces from it and incorporate them into your own approach. I am a software-tools kind of guy. There are probably numerous IDEs for Web development that will do everything for you once you know them, and there are likely a number of Eclipse plugins. Powerful, dedicated tools typically have a steep learning curve that pays off only if you do a lot of that type of work.

This article is not about DocBook XML. It is about how to build Web sites using CSS to render DocBook XML documents simply. I am not a Web developer, and I opt to learn tools that have broad uses. The tools I use for building Web content are vim for editing, m4 or Perl for macro processing and HTML tidy for verification—the same tools I use to develop software and write documentation. During the past few years, I have added basic XML, particularly DocBook XML, to my list of fundamentals.

I keep a simple DocBook XML article template readily available and pull it up in vim whenever I feel inspired to write something technical that is larger than an e-mail. By using a DocBook XML template, I can focus mostly on content and produce results that are clear and meaningful, with minimal emphasis on presentation.

More recently, I have discovered that with a little help from CSS, DocBook XML documents can be viewed directly on any Web site by CSS-capable browsers, without transforming to HTML, making it easy to add to my Web site. For more complex documents, OpenOffice.org supports DocBook XML as an output format, and there are increasingly more tools to produce and manipulate DocBook XML. DocBook XML can be read directly by OpenOffice.org or transformed easily into all commonly used document formats, such as HTML, PDF, Word and so on. One objective of XML (one that would be difficult to identify in the competing XML word-processor formats) is divorcing content from presentation. This is a principle I heartily endorse.

I make a distinction between the parts of a Web site used for navigation and the content of the Web site. I deliberately choose to separate content physically from navigation. With rare exceptions, all content pages are devoid of navigation and function as standalone documents. Today, I do them in DocBook XML. Previously, I used HTML; however, I always tried to maintain a separation between content and navigation. My first step is to build an HTML presentation/navigation framework. I create the main HTML index page for the site, and I use HTML FRAMES to divide the display into three regions: a header, a menu and a body. FRAMES are somewhat frowned upon within Web development, as they can be used to capture other people's Web content and create the impression that it is your own. They also can impede navigation, and they may be less friendly to people with disabilities. However, I am not aware of another equally easy-to-use Web construct that can be made to separate content from navigation and presentation. There are other means to achieve similar effects, but all of those that I am aware of incorporate navigation and presentation elements into the content. My objective is to be able to develop the content of the Web site in DocBook XML, modified only to include a stylesheet and to isolate presentation and navigation elsewhere.

There is one other heretical side effect to this approach—nothing about it requires a Web server. You can build and test all of this in the browser of your choice without installing a Web server, and when finished, you can drop it all on a CD-ROM where it can be viewed on any system with a Web browser.

The core of my index page is:

```
<frameset class="frame" cols="140,*" bordercolor="#000000"
frameborder="0" framespacing="0">
  <frame class="frame" src="margin.html" name="Margin" scrolling="no"
marginwidth="0" marginheight="0"
  <frameset class="frame" rows="100,*" bordercolor="#000000"
frameborder="0" framespacing="0">
  <frame class="frame" src="header.html" name="Header" scrolling="no"
marginwidth="0" marginheight="0" />
  <frame class="frame" src="home/index.xml" name="Body" scrolling="auto"
marginwidth="0" marginheight="0" frameborder="0" />
</frameset>
</frameset>
```

This divides the browser display into three regions. A menu area on the left, a header at the top and a body for content in most of the remainder. My header page tends to be fairly trivial, basically:

```
<body class="header" id="body-header">
  <div class="header">
    <h1 class="header">My Title</h1>
  </div>
</body>
```

The class and id tags allow the use of CSS to overload style later. The margin is almost as simple:

```
<body class="margin" id="body-margin">
  <div class="menu-box">
    <div class="menu" id="home">
      <a href="home/index.xml" target="Body">Home</a>
    </div>
    ...
  </div>
</body>
```

Again, the class and id tags are for CSS style. The menu-box block element surrounds all the menu items. The menu block elements can be repeated as needed. CSS can be used to style the menu items to suit personal taste. Specifying a target for the links means that when a menu item is clicked on, it changes the document in the "Body" frame of the frameset.

I use the following CSS to create highlighted menu buttons:

```
div.menu-box {
  display: block;
  border-width: 2pt;
  border-color: color_bkgr !important;
  border-style: inset ;
}

div.menu {
  border-style: inset ;
  border-width: 5px ;
  background: color_menu_bkgr1 !important;
  border-color: color_menu_bkgr !important;
  color: color_bkgr !important ;
  font-weight: bold;
  font-size: 8pt;
  height: 14pt ;
  width: 110pt;
  vertical-align: middle;
  x-margin: 5pt;
  x-padding: 5pt;
  text-align: center;
  padding-left: 5pt;
}

div.menu:hover {
  position: relative;
  top: 1px;
  left: 1px;
  border-color: color_menu_bkgr1;
  background-color: color_menu_bkgr;
}

a.menu { text-decoration: none }
```

Those are all the key elements of the non-content portion.

The menu system can be nested. Changing the target of a menu item to "Margin" can pull in a new side menu, and that can be repeated as often as you like. Internet Explorer's handling of CSS, particularly positioning, is broken, so there are subtle differences in

the display between it and properly conforming browsers. Complicated cross-browser CSS positioning can be extremely difficult, and it is further complicated because Internet Explorer 7 is slated to fix many CSS issues in ways that break most of the published work-arounds for earlier versions. Also, I would advise being careful about background colors. I spent a short life time failing to figure out how to eliminate a white streak between the menu area and the body that appeared only with Internet Explorer and only if I used a background color. This article is not about how to become proficient at fancy cross-browser Web development; the focus is on providing a simple approach to easily display content that looks pleasant, regardless of the browser. Getting pixel-for-pixel identical CSS cross-browser results for numerous browsers is a complex task.

Up to this point, I have ignored the HTML headers and issues, such as the fact that `color_menu_bkgr` is not a valid HTML/CSS color.

HTML pages, such as `index.html`, `header.html` and `margin.html` need valid HTML headers, and they need a link element referencing the CSS stylesheet, such as:

```
<link rel="stylesheet" type="text/css" href="/css/stylesheet.css"
title="default">
```

added to the header.

The CSS excerpt above is from `stylesheet.css`, which also can include any additional CSS you might want to add or overrides for the default DocBook CSS. A number of CSS stylesheets are available for DocBook XML—several are listed on the DocBook Wiki, and the particular stylesheet I use is `badgers-in-foil` (see the on-line Resources). The `badgers-in-foil` stylesheet has allowed me to render DocBook XML articles pleasingly in several different browsers.

All XML pages need two stylesheet links added to the XML header:

```
<?xml-stylesheet href="/css/docbook-css/driver.css" type="text/css"?>
<?xml-stylesheet href="/css/stylesheet.css" type="text/css"?>
```

The second link is not strictly necessary, but it can be used to override or add additional style information to the DocBook XML files, without changing the DocBook XML stylesheet.

I handle the generation of the framework, XML and HTML wrappers and many repeated elements using the macro processor `m4`. It could be done as easily with Perl or `bash/sed`. This allows me to define standard headers, colors and other useful string substitutions as `m4` macros. `color_bkgr` is an `m4` macro and will be replaced by `m4` with the background color I have chosen for this site anywhere it occurs. I reuse the same framework whenever I need to create a new Web site. I can create a new site with different content, titles, colors and so on by changing a few macros. However, the complexity gradually has increased to the point where I am starting to think of moving from `m4` to Perl for the preprocessing. I am using automated generation of XML and HTML, and therefore it is an excellent idea to use HTML `tidy` after processing to verify it.

First, install HTML `tidy` and `m4`. I primarily work with Debian and Debian derivatives, so installing `tidy` and `m4` consists of:

```
apt-get install tidy
apt-get install m4
```

Most distributions should provide `m4` and have `tidy` available through their package system. See Resources for the main pages for `tidy` and `m4`.

Then, I have a text file (`pages.list`) with a list of the base names for all pages, as well as their type: CSS, HTML and XML:

```
stylesheet.css
index.html
header.html
margin.html
home.xml
...
```

I use a short shell script to run `m4` and HTML `tidy` on each page and place the results where they belong:

```
#!/bin/sh
# $Id:
# $URL:

#dest=../test
dest=..
lname=pages.list

dopage() {
  echo "$1"
  if [ "$2x" == "xmlx" ]; then
    if ! [ -d $dest/$1 ]; then
      mkdir $dest/$1
    fi
    m4 -D_xml $1.m4 | tidy -i -xml >$dest/$1/index.xml
  elif [ "$2x" == "htmlx" ]; then
    m4 $1.m4 | tidy -i >$dest/$1.html
  elif [ "$2x" == "cssx" ]; then
    m4 -D_css $1.m4 >/var/www/share/css/$1.css
  else
    echo "Whoops $1 $2"
  fi
}

if [ -f $lname ]; then
  list='cat $lname | grep -v '#' | awk '{print $1}' | tr '\n' ' '`
  for argv in $list ; do
    page="" ; fmt=""
    page=`echo $argv | awk -F "," '{print $1}'`
    fmt=`echo $argv | awk -F "," '{print $2}'`
    dopage ${page} ${fmt}
  done
fi
```

Now, `m4` can handle the generation of standard headers, links to stylesheets, macro substitutions, substitutions for color names and so forth. The menu items even can be generated automatically from macro items.

The `header.m4` file to generate the header page becomes:

```
define(_page.header)dnl
include(defs.m4)dnl
include(hdr.m4)dnl
<div class="header">
  <h1 class="header">_title</h1>
</div>
include(ftr.m4)dnl
```

A Web server is not needed to view any of the framework and content we have created, but most Web pages are distributed by a Web server. No additional configuration should be needed for most Web servers; however, the following CSS config file added to `/etc/apache2/conf.d` creates an alias, allowing the CSS directory to be shared across multiple sites or to be referenced easily regardless of the relative path inside the Web site:

```
Alias /css /var/www/share/css/

<Location /css>
  Order allow,deny
  Allow from all
  Options Indexes FollowSymLinks MultiViews
</Location>
```

This is a software-tools approach. For a small number of Web sites with very little content, there is no benefit to adding the complexity of automating

the generation of HTML or XML headers and footers. Where there is a significant amount of content, frequent modification or numerous unique sites, there can be a substantial benefit.

I have barely touched on DocBook XML. I started "word processing" in college using text formatters like runoff, nroff and text on my H8. The concept of separating content from appearance is a natural return to my non-WYSIWYG word-processing roots.

There are tools available to do WYSIWYG processing of XML documents. The easiest approach, if you are more comfortable with a WYSIWYG word processor, is to use OpenOffice.org, which can save documents as DocBook XML. OpenOffice.org's DocBook XML capabilities are limited, however. It is not typically possible to go from a well-formatted OpenOffice.org format or Word format file to a DocBook XML document without losing some facets of the presentation. Plain DocBook XML is more focused on content and structure than presentation details. OpenOffice.org does not associate a stylesheet with the saved DocBook XML document, so style items, such as typefaces, type size, indents and so on, will be supplied by the DocBook XML CSS you use. If you are not completely happy, you either can modify the stylesheet or override it by "cascading" a new stylesheet, changing the elements you want to change.

As I mentioned previously, I am happy with the badgers-in-foil stylesheet. My CSS makes very few changes. I am more focused on creating readable documents easily and getting them to my Web site or transforming them into other file formats as needed. As I mentioned, I usually choose to start with a simple DocBook XML article template. I use vim to add my content to that template. The template uses a bare minimum of DocBook XML, and aside from some XML fundamentals, such as making certain that start and end tags remain matched, my paragraphs use little more than a few very obvious tags.

Proficient DocBook XML users can master a rich set of DocBook XML constructs, but ordinary users can easily produce increasingly sophisticated documents by slowly learning only a few tags. I find DocBook XML significantly easier to use than HTML. XML is rigid in tag matching and nesting rules, and there are less, if any, idiosyncrasies. Structure and organization—lists, tables, paragraphs, chapters, sections and so on—are all done in DocBook XML. Appearance and presentation decisions are made in the stylesheets. Capable CSS developers could transform a basic DocBook XML article into something elegant. However, my objective is not elegant documents and Web sites, but in making content informative and readable in a variety of formats quickly and simply.

DocBook XML is an increasingly popular approach to constructing Web documents. Numerous open-source projects, as well as the Linux kernel, are relying more heavily on DocBook XML as a standard format for documentation. The Linux Documentation Project provides an author's guide with the sample article template I frequently use, as well as a large number of links to other DocBook XML resources. Eric Raymond's "DocBook Demystification HOWTO" provides an excellent explanation of why DocBook XML is important and why it is replacing most other formats for open-source documentation. Michael Smith's "Take My Advice: Don't Learn XML" is similar and explains why making worthwhile use of DocBook XML does not have to involve becoming an expert in XML or the plethora of associated XML technologies. *The Definitive Guide* by Norman Walsh and Leonard Mueller will provide you with much more than you likely will need to know, as well as critical answers if your use of DocBook XML starts to become more sophisticated. And finally, I hope this article makes clear that making effective use of DocBook XML can be simple and requires developing minimal new skills. ■

Resources for this article: www.linuxjournal.com/article/9263.

Dave Lynch is a software consultant. Web development, XML, CSS and HTML are occasional tangential elements of the embedded and systems software that he writes, usually under Linux, in a vain attempt to make a living. In another life, he is an architect, and he currently keeps himself occupied when not wreaking havoc with his Web site or writing software for clients by building his own home.



Great Minds Great Solutions

1-877-25-SERVER www.genstor.com

Customized Solutions for

SERVERS :: **STORAGE** :: **APPLIANCES**

Linux - FreeBSD - x86 Solaris - MS etc.

SERVERS

LOW POWER - BUDGET - HIGH DENSITY

(1U, 2U, 3U and above - Xeon, P4, Opteron, VIA)



2U Dual Xeon/Opteron Server
Upto 24GB RAM
Upto 8 SATA/SCSI HDD
Starting @ \$ 2000.00
Call for your custom needs

SERVER MANAGEMENT & MORE

Genstor offers Data Center-proven Linux Management Technology - which gets you :



Provisioning
Deployment
Change Management for Linux
Environments large and small
Track Changes
Instant Rollback and much more

Setup your racks with heterogeneous Linux environments in minutes
Contact sales@genstor.com

STORAGE

SATA- NAS-DAS-iSCSI -SAS Storage Solutions



5U Dual Xeon/
Opteron SATA Storage
Upto 24GB RAM
Upto 26 SATA HDD
Upto 18TB Storage
Starting @ \$ 3930.00
Call for your custom needs

CUSTOM APPLIANCE SOLUTIONS

Prototype - Certifications - Logo Screen Printing



Custom Turnkey OEM Appliance Solutions
From Prototype to Drop Ship
Custom OS/Software Image Installs
No Quantity is small to Customize
Call for your Custom OEM Solutions

FACING POWER PROBLEMS AT DATA CENTER PL. CALL FOR SOLUTIONS

Contact Genstor For Your Hardware Needs

Genstor specializes in customizing hardware around the OS of your choice (Linux, *BSD, x86 Solaris etc.) with Intel and AMD offerings. Please contact Genstor sales@genstor.com for all your hardware needs.

GENSTOR SYSTEMS, INC.

780 Montague Exp. # 604, San Jose, CA 95131

Phone: 1-877-25 SERVER or 1-408-383-0120

Email: sales@genstor.com

Fax: 1-408-383-0121

Prices, Products and Availability subject to change without notice.

Linux and Open Source in Telecommunications

A success story with a focus on Carrier Grade Linux. IBRAHIM HADDAD

This article examines the impact that Linux and open-source software are having on the telecommunication industry, technology trends moving toward open and standards-based platforms and the .orgs that are active in promoting carrier grade base platforms. Furthermore, this article focuses on the Carrier Grade Linux initiative at OSDL and discusses its contributions to this growing ecosystem.

Introduction

The telecommunication industry is facing several challenges:

- Telecom service providers are looking to reduce their costs using commodity software and commercial off-the-shelf (COTS) hardware building blocks.
- Telecom service providers require seamless integration of COTS carrier grade components; the integrated solution must be validated for carrier grade availability.
- The growth of packet traffic is putting pressure on communication networks originally designed for "store and forward"; platforms in an all-IP environment that maintain carrier-class characteristics are delivering increasing levels of availability and dependability.
- Operators want to decrease time to market and increase the capability for fast delivery of new services by shortening new service development time and unifying platforms.
- Of course, operators want to roll out the above capabilities while still making money and increasing profits.

Linux and open-source software provide a compelling avenue to operator success. The open-source operating system has certain characteristics that confer upon it advantages over other operating systems; indeed, Linux has been a disruptive technology with clear impact in telecommunications. Today, not only do many of the server nodes with telecom networks run Linux, but Linux also powers mobile phones and many intermediate nodes "in the middle".

So, what is a disruptive technology, and how does it impact an industry?

Disruptive Technology

Disruptive technologies first appear saddled with significant deficiencies and are usually targeted at niche segments. Disruptive technologies, however, also provide significant cost benefits. For example, a truly disruptive technology may offer only half the performance of its legacy competitor but can be delivered at one-tenth the cost.

Disruptive technologies are most often taken up by early adopters and then experience a much slower adoption into

the mainstream. The adoption of a disruptive technology always starts with non-mission-critical applications (such as utility computing) and moves to mission-critical application as it matures (such as business-critical and enterprise core applications). Linux adoption followed this pattern, starting out hosting Web servers, e-mail and FTP servers and moving now to mission-critical applications, such as telephony. With increasing adoption, a disruptive technology, such as Linux, provides an opportunity (or even forces) companies to re-evaluate and also re-invent their business models and identify real value-added products and services. Companies that do not provide clear value quickly find themselves out of the market.

Linux adoption in telecommunication has not only been increasing, but adoption is also accelerating. Reasons to adopt Linux vary but revolve around common key advantages, such as licensing terms, full access to source code, freedom to choose from multiple providers, lower costs versus legacy and proprietary operating systems, higher system performance, reliability, security, source code quality, innovation rate, peer review, testing resources and the availability of an established ecosystem.

The Telephony Business in (R)Evolution

The traditional telecommunication business model is one of high-margin and high-revenue business. In the past, telecom experienced better than 10% year-on-year growth, and almost any project could become successful because demand was so great. Telecommunication companies bought in and sold on proprietary solutions, taking a margin on top of the initial licensing costs. Standards were sufficient only to ensure basic connectivity; after that, essentially proprietary models were built up, with vendor lock-in as the norm.

Figure 1 illustrates the state of the telecom business beginning in the mid-1980s up to the present. In the 1980s, the carrier's business was monopoly-based with very few players in the field, which provided carriers the opportunity to make a lot of money, due to significant margins with voice telephony as a high-priced premium service. In the mid-1990s, new players (carriers/operators) entered the business, increasing the competition. However, voice telephony was still a

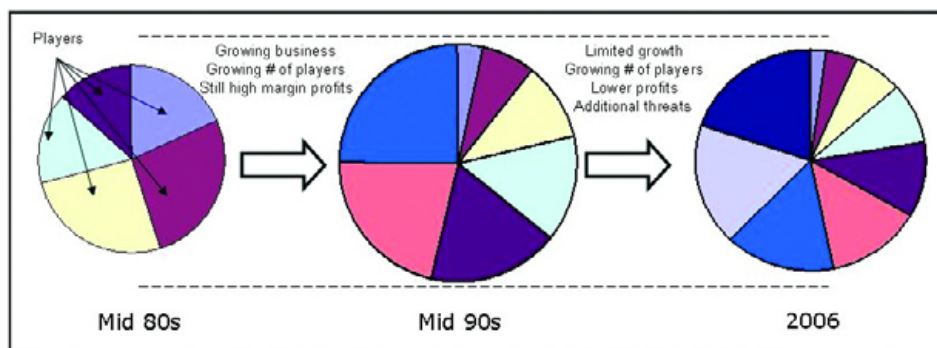


Figure 1. Telephony Business in (R)Evolution: the Business State from the mid-1980s to Today

premium service, and although prices were falling, operators still had significant margins. Today, the business looks very different. It is shrinking with many more players in the space, increased competition and much diminished profits. Voice telephony is a commodity. Furthermore, the industry faces additional threats, such as VoIP and broadband telephony. How can they beat free or close-to-free calling?

Technology Trend

In the past (circa 1985), communications and data service networks built on proprietary platforms to meet specific requirements for availability, reliability, performance and service response time. However, communications service providers needed to drive down costs while maintaining carrier-class platforms with high availability, scalability, security, reliability, predictable performance and easy maintenance and upgrade.

The current technological trend in this space, illustrated in Figure 2, is moving away from expensive proprietary and legacy systems consisting of proprietary technologies and components without a clear separation of the "building blocks" into standards-based systems that consist of interchangeable software and hardware COTS "building blocks" that communicate with each other using standardized interfaces and that are offered by multiple providers.

Traditionally, communications and data service networks were built on proprietary platforms that had to meet very specific requirements in areas such as availability, reliability, performance and service response time. Those proprietary systems were composed of highly purposed hardware, operating system and middleware and often included proprietary technologies and interfaces. Such proprietary approaches to system architecture fostered vendor lock-in, served to limit design flexibility and freedom and produced platforms that are very expensive to maintain and expand.

Today, those same service providers and carriers are challenged to drive down costs while still maintaining carrier-class characteristics for platforms to provide service and mission-critical applications in an all-IP environment. Providers are in a position today where they must move away from specialized proprietary architectures and toward COTS approaches and building practices (Figure 2) for several reasons:

1. Faster time to market.
2. Reduced design and operation costs by using COTS hardware and software components.
3. The growth of packet traffic is placing added pressure on communication networks. Communication platforms reside on all-IP networks and need to maintain carrier grade characteristics in terms of availability, reliability, security and service response time.
4. The emergence of COTS hardware and software components is driving the need for seamless integration of all components as integrated solutions that must be

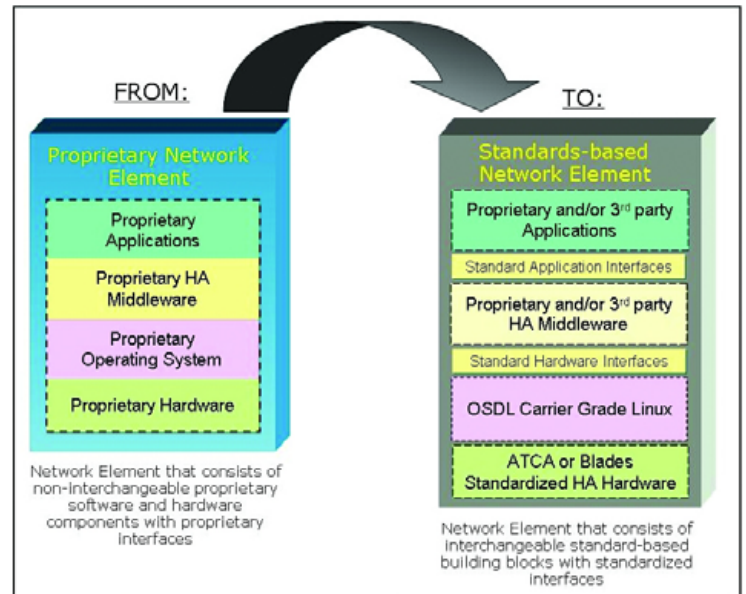


Figure 2. The Technology Trend from Closed Proprietary to Open Standards-Based Platforms

DEDICATED SERVERS ROCK BOTTOM PRICING



\$60 /mo
Celeron D

\$160 /mo
Dual Core Xeon

- » 1300GB/mo Bandwidth
- » 1GB RAM DDR400!
- » 160 GB SATA2 Hard Drive
- » 100Mbps Dedicated Port
- » FREE REBOOTS 24x7
- » FREE SUPPORT 24x7
- » FREE SETUP Available*
- » BRAND NAME Hardware

w.cari.net

* See website for details.

carinet
betterServers. betterService.

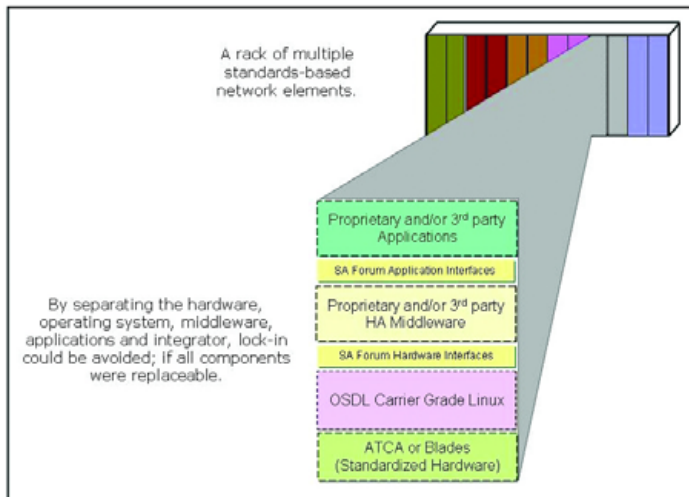


Figure 3. A Typical Telecom Rack with Multiple Network Elements

validated for carrier grade availability and scalability.

The benefits of a standardized platform (Figure 3) based on COTS hardware and software are many:

1. Avoiding lock-in: by separating the hardware, operating system, middleware, applications and integration, vendor lock-in can be avoided by making components replaceable and interoperable through standardized interfaces.
2. The platform achieves economic as well as technical scaling.
3. All components and ecosystem links, including the integrator, can be changed if they underperform, with minimal impact.
4. A fully open-source route is possible for next-generation networks and products.
5. End customers benefit from multiple products running side by side on the platform and from an improved cost base and speed from fewer adopted platforms.
6. Moving to CGL from a proprietary OS can save telecom equipment manufacturers money because they don't have to develop, maintain or license an in-house proprietary OS. Instead, they can invest in the CGL ecosystem to make Linux good for their own use. In addition, the flexibility of an open-source operating system provides for more customization, increasing each manufacturer's competitive advantage.

To summarize, the telecommunication industry is transitioning to COTS architectures and practices, embracing Linux and open-source software and re-aligning at multiple levels. Before

Vendor Lock-in

Lock-in is an economic issue, not a technical one. It presents a technology "exit barrier" and takes four steps. First, vendors' offers initially vary—with low cost but proprietary solutions, well-integrated by having just enough standard interfaces and APIs (proprietary is often called "value added" or something similar). Next, vendors offer business-case compelling information, based around the presumed low cost of their solution. The

1999/2000, the industry experienced incompatible platforms, protocols, high barriers to entry, circuit switches and so on. Today, the telecommunication industry is resurging with COTS, Linux and open-source software, with many new players and many opportunities for new businesses.

The .org Players

There are five major .orgs (Figure 4) active in the space of accelerating the adoption of carrier grade platforms that are based on COTS hardware and software. These organizations are CP-TA, OSDL, PICMG, SA Forum and the SCOPE Alliance. In the following sections, we present each of these organizations, discuss their goals and highlight their contributions.

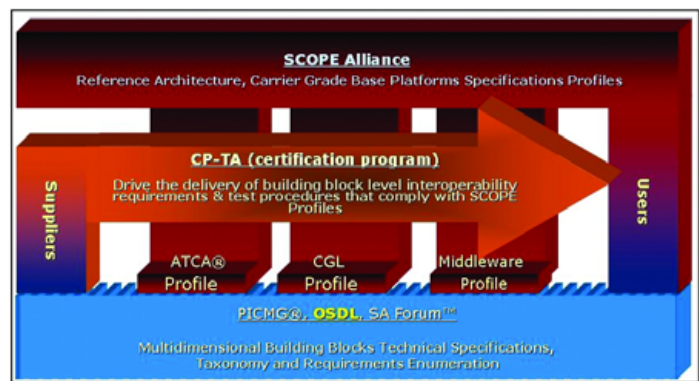


Figure 4. The .org Players: PICMG, OSDL, SA Forum, CP-TA and the SCOPE Alliance (courtesy CP-TA)

CP-TA

The Communications Platforms Trade Association (CP-TA) is a group of communications platform and building block providers dedicated to accelerating the adoption of SIG-governed, open specification-based communications platforms through interoperability testing and certification. With industry collaboration, CP-TA plans to drive a mainstream market for open industry standards-based communications platforms.

SA Forum

The Service Availability Forum (SA Forum) is a consortium of communications and computing companies working together to develop and publish high-availability and management software interface specifications.

OSDL CGL

The OSDL Carrier Grade Linux (CGL) initiative is an industry forum that supports and accelerates the development of Linux functionality for

third step is encouraging a strong roll-out of the solution to establish a sufficiently large installed base to start raising costs (license, support and so forth). The fourth and final step is when suppliers raise pricing up to, but not beyond, the point where additional roll-out of their equipment is slightly less expensive than replacing everything with an alternative vendor's. The exit barrier has been raised, and you are now locked in.

telecommunication applications. The goal of CGL is to make Linux better for the telecommunication industry. A Linux kernel with carrier grade characteristics is an essential component in open, standards-based communication platforms and architectures. OSDL specifically focuses its work on the Linux operating system and collaborates with other industry organizations to drive adoption of open standards and open-source software. It works closely with each group to ensure that efforts are complementary and deliver value to the market.

SCOPE Alliance

The SCOPE Alliance is an industry alliance committed to accelerating the deployment of carrier grade base platforms for service provider applications. Its mission is to help, enable and promote the availability of open carrier grade platforms based on (COTS) hardware and software and Free and Open-Source Software (FOSS) building blocks and to promote interoperability to better serve service providers and consumers.

PICMG

The PCI Industrial Computer Manufacturers Group (PICMG) is a consortium of more than 450 companies who collaboratively develop open specifications for high-performance telecommunications and industrial computing applications. The consortium has resulted in a series of specifications that include CompactPCI, AdvancedTCA, AdvancedMC, CompactPCI Express, COM Express and SHB Express. The goal of PICMG is to offer equipment vendors common specifications, thereby increasing availability and reducing costs and time to market.

Carrier Grade Linux Initiative at OSDL

The OSDL Carrier Grade Linux working group was established in January 2002. Its goal is to identify requirements for enhancing the Linux operating system to achieve an open-source platform that is highly available, reliable, secure and scalable, and suitable for carrier grade systems. The CGL working group has the vision that next-generation and multimedia communication services can be delivered using Linux-based platforms. To realize this vision, the work group developed a strategy to define the requirements and architecture for the Carrier Grade Linux platform and promote development of a stable platform for deployment of commercial components and services.

The CGL working group focuses on two areas: carrier grade enhance-

CGL and the COTS Ecosystem

CGL is an important part of the telecommunications move to using COTS components for building equipment. Carriers and service providers are in a position today where they must move away from specialized proprietary architectures toward COTS approaches and building practices for several reasons, such as reducing design and operation costs and gaining the ability to deliver new services faster based on common standardized platforms. In addition, the increased power and reliability of such building blocks, along with the development of specifications that guide their implementation, are allowing more flexibility for network deployment with improved price performance. CGL is a core building block, providing a Linux kernel that offers the needed reliability, availability and performance for platforms running in mission-critical environments and providing communication services.

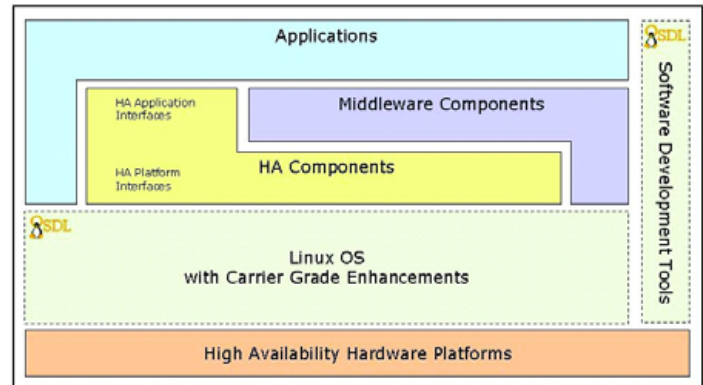


Figure 5. Scope of the Carrier Grade Linux Working Group

ments to the operating system that are related to various requirements, such as availability and scalability, and software development tools. Today, more than two-dozen OSDL member companies from all over the globe are actively involved with the CGL initiative. Member companies cover the whole ecosystem: carriers, network equipment providers (NEPs), telecom equipment manufacturers (TEMs), platform providers, independent software vendors (ISVs), middleware providers and Linux distributors.

The CGL working group also identifies existing open-source projects that map to the CGL requirements. The result is the CGL Development Guideline Web site (see the on-line Resources). This is an effort from the CGL initiative to survey open source for projects that can potentially provide implementations for the requirements defined

The advertisement features a black and white photo of Linus Torvalds, the creator of Linux, with his characteristic long hair and beard, wearing sunglasses and making a peace sign. The text above the photo reads 'Linux Laptops Starting at \$799' and 'DON'T BE SQUARE! GET CUBED!'. At the bottom, there is a logo for '3 Technologies' with the tagline 'Working hard to bring Technology to Life', and the contact information '309.34.CUBED shoprcubed.com'.

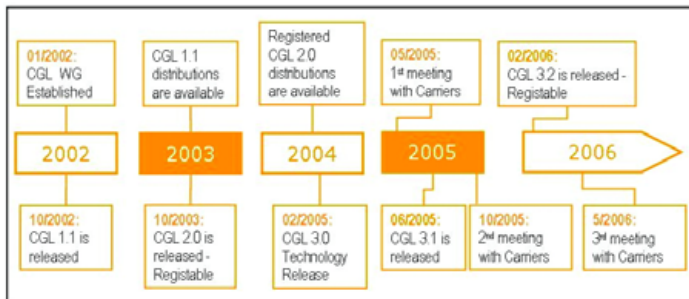


Figure 6. Overview of the CGL Initiative from Its Inception in 2002 to June 2006

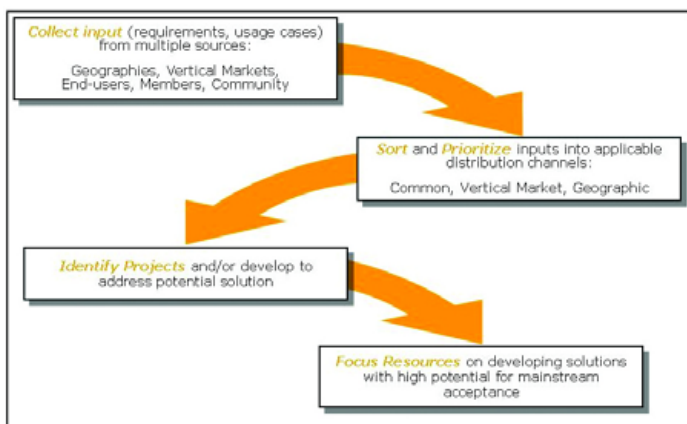


Figure 7. The CGL Initiative Work Process

in the CGL Requirements Documents. This site is maintained and updated frequently.

The CGL working group collects requirements from multiple industry sources and develops use cases for the various proposed capabilities and functionalities. The working group then sorts and prioritizes the input from the industry, member companies and end users to identify open-source projects that are working on these areas. If no open-source project exists, the working group starts new open-source projects to develop these capabilities and focuses its resources to develop solutions with high potential for mainstream acceptance. In many instances, member companies have (re)released previously proprietary technologies as open source to accelerate the availability of these capabilities in Linux.

The CGL initiative released the original CGL Requirement Definition Document in 2002 (v1.1) and has issued two revisions (v2.0 and v3.2), and it also has established a registration process for Linux vendors to register compliance of their Linux distributions.

Debian, the Latest Compliant Distribution to CGL

The CGL working group established a registration process for Linux distributions to disclose information on how they meet the CGL requirements. The process is a public disclosure of all CGL requirements as mandated by each CGL release version and describes how the Linux vendor met the CGL requirements. The outcome of the registration process allows CGL-registered platform suppliers to market their Linux distributions and systems to NEPs and TEMs and carriers with the CGL registration mark to demonstrate the platform's suitability for carrier grade applications.

In June 2006, Debian passed the CGL 2.0 registration process, becoming the seventh distribution that meets the CGL 2.0 requirements. The other six are Asianux, FSMLabs, MontaVista, Novell, TimeSys and Wind River. The Debian announcement is of great importance. Debian is one of the leading distributions of the Linux operating system. Its registration adds more than 1,000 developers and tens of thousands of end users to the CGL community. Debian registration gives telecommunications providers a fully open platform that comes with the support of one of the strongest Linux communities and represents an ideal balance between "roll-your-own" CGL solutions and available commercial options. Telecommunications equipment providers looking for a fully open option now have one.

Closing

In the February 2006 *LinuxWorld* magazine editorial, "The Holy Grail of Networking", Stuart Cohen, CEO of OSDL, discussed the end-to-end infrastructure with a single operating system (Linux) and the role OSDL is playing to enable this single OS infrastructure from the server to the handset. At OSDL, the CGL and MLI initiatives are driving forward an "end-to-end" Linux deployment, succeeding in its mission to accelerate the development and adoption of Linux from the enterprise to mobile computing in a vertical industry that has been historically dominated by proprietary technologies. What's next for Linux? Only time will tell.

To learn more about how OSDL initiatives are helping accelerate the development and adoption of Linux, visit the OSDL Web site (see Resources).

Acknowledgements

The author would like to thank Bill Weinberg, OSDL's Senior Technology Analyst, for his valuable reviews and contributions. ■

Resources for this article: www.linuxjournal.com/article/9267.

Ibrahim Haddad manages the Carrier Grade Linux and the Mobile Linux Initiatives at OSDL, promotes the development and adoption of Linux in the Communication industry, and leads the Carriers/NEPs Forum he established at OSDL in early 2005. Prior to joining OSDL, Ibrahim was a Senior Researcher at the "Research and Innovation" Department of Ericsson Corporate Unit of Research in Montréal, Canada, where he was involved with the server system architecture for 3G wireless IP networks and contributed to Ericsson's open platform efforts. Ibrahim is co-author of two books on Red Hat Linux and Fedora, and a Contributing Editor of three leading Linux publications. Ibrahim received his PhD in Computer Science from Concordia University in Montréal, Canada.

CGL Initiative Achievements:

- Increasing the number of OSDL member companies involved with CGL; the latest members include Siemens and Motorola.
- Three major releases of the CGL Requirement Definition Documents: CGL V1.1 in October 2003, CGL 2.0 in October 2003 and CGL 3.2 in February 2006.
- Seven distributions and Linux vendors registered for CGL 2.0:

TimeSys, Novell/SUSE, MontaVista, FSMLabs, WindRiver, Asianux and Debian. Linux vendors are now in the process of registering for compliance with CGL 3.2.

- More than 25 platform providers are integrating CGL as part of their offering.
- Service providers and carriers are deploying CGL-based platforms.

Growing a World of Linux Professionals



We at the Linux Professional Institute believe the best way to spread the adoption of Linux and Open Source software is to grow a world wide supply of talented, qualified and accredited IT professionals.

We realize the importance of providing a global standard of measurement. To assist in this effort, we are launching a Regional Enablement Initiative to ensure we understand, nurture and support the needs of the enterprise, governments, educational institutions and individual contributors around the globe.

We can only achieve this through a network of local "on the ground" partner organizations. Partners who know the sector and understand the needs of the IT work force. Through this active policy of Regional Enablement we are seeking local partners and assisting them in their efforts to promote Linux and Open Source professionalism.

We encourage you to contact our new regional partners listed above.

Together we are growing a world of Linux Professionals.



Stable. Innovative. Growing.

SMART (Smart Monitoring and Rebooting Tool)

If you want an agent to monitor and control services, you'll need to get SMART.

ALBERT MARTORELL

There are a lot of excellent monitoring tools (Big Brother, Nagios and so on), and some of them allow recovery from dead services, but with great complexity in their configuration, which becomes even more complicated when you want to supervise local services that are not remotely accessible, such as syslog, xinet, mrtg, iptables or Nagios itself.

The purpose with SMART was to have a simple, flexible and quick-to-implement application for monitoring the most critical system daemons that made it possible to add new ones without modifying the code and to avoid installation and configuration complexities. It also needed to be capable of making decisions and solving problems (or at least trying to do that).

Evolution

After a first version of "passive" monitoring, we tried to go a step further and obtain an "active" application, that is to say, to add the possibility of auto-recovery. By executing the application periodically through cron, it should detect daemons that were down and boot them without the intervention of the system administrator.

Later, we considered the possibility that a nonprivileged user could execute this application from a console or remotely (via Telnet or SSH). Centralization of detection and error recovery in only one script made integration with sudo easier. Furthermore, it allowed delegating some stronger recovery actions needed in critical situations, such as rebooting the whole system, to this nonroot-privileged user.

With the ps command, we can list all the active processes in the system, but being "active" is not the same as being "operative", so this led us to include the check scripts, which are small programs to test services and determine whether they really are operative and answering requests. The difficulties we found suggested that we not waste efforts re-inventing the wheel and profit from plugins included in Nagios (monitoring software that we were using satisfactorily for almost three years).

Files and Directories

The distribution of SMART has two shell scripts (smart and check-service), two configuration files (host.conf and services.conf) and two directories (scripts and plugins), which contain the check scripts and the plugins (Listing 1).

Permissions of files and directories allow a nonprivileged user called sysman to execute the application, but deny sysman the ability to modify the contents to use it in an inadequate way.

General Operation

The SMART program reads the configuration files services.conf and host.conf and executes check-service

Listing 1.

The SMART Installation Files and Directories

```
[root@server ~]# ls -la /home/sysman/
drwxr-x--- 4 root sysman 4096 May 27 11:49 .
drwxr-xr-x 3 root root    4096 Jul  8 2003 ..
-rwxr-x--- 1 root sysman 1448 May 27 11:51 smart
-rwx----- 1 root root    7815 May 27 11:51 check-service
-rw-r--r-- 1 root root     242 May 27 11:49 host.conf
drwx----- 2 root root    4096 Apr 29 13:38 plugins
drwx----- 2 root root    4096 Apr 29 13:39 scripts
-rw-r--r-- 1 root root     883 May 17 10:40 services.conf
```

for each defined service. If a check script has been assigned to a service, for example, services 1 and 2 in Figure 1, check-service will execute it, passing the needed parameters and then will wait for the exit status to determine whether the service is alive. If this check script executes some other external script (plugin), such as service 1 in Figure 1, this one will be responsible for checking the service status.

If no check script has been assigned to a service (service 3 in

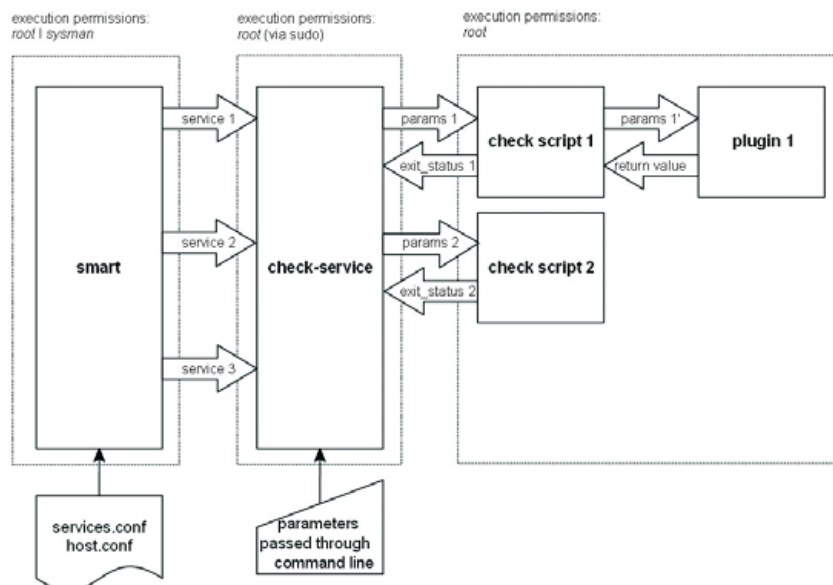


Figure 1. SMART Program

Figure 1), the check-service file will determine the service status by getting the number of active processes. According to this information, the SMART command-line parameters and the configuration parameters, it will decide what actions to carry out.

Integration with sudo

Integration with the sudo (superuser do) tool allows the system administrator to permit another user (sysman) to start dead services, restart all the services or reboot the whole system. Advantages of this are:

- Simple configuration: there's no need to give privileges to that user to stop and start every service, and no need to use administrative tools (ps, kill, rm and so on). The check-service script centralizes the whole operation.
- Security: user sysman can't read, write or execute the check-service file.
- Easy to use: scripts are managed by sudo, so its usage will be transparent for the user.

For a user sysman, who needs privileges on the host server, the configuration file of sudo (/etc/sudoers) should be as shown in Listing 2.

Listing 2.

sudo Configured for SMART Access

```
# Defaults specification
Defaults:root !syslog

# User privilege specification
root    ALL=(ALL) ALL
sysman  server=(root) NOPASSWD: /home/sysman/check-service
sysman  server=(root) NOPASSWD: /sbin/reboot
```

This way, we disable syslog logging when sudo is executed by user root, and we assign root privileges to user sysman, at the host server, only for the execution of commands /home/sysman/check-service and /sbin/reboot, without asking sysman for the password every time.

Verifications

Through the PID file defined in the configuration file, we obtain the parent process identifier (PID), and we determine the number of active processes generated by this service. Next we check whether:

- The service is responding to petitions within the defined time period.
- The number of processes generated by the service doesn't exceed the maximum and minimum defined thresholds.

Status Determination

Considering the results obtained in former verifications, we classify the service status:

- 0: service is responding to requests within the defined time period, the number of processes generated by service remains between the defined thresholds, and the information provided by the PID file is correct.

- 1: service is responding to requests within the defined time period and the number of processes generated by service remains between the defined thresholds, but either the information provided by the PID file is incorrect or this file doesn't exist, even though it has been defined.
- 2: service is responding to requests within the defined time period, but the number of processes generated by the service is beyond the defined thresholds (this could be the case of an overloaded but operative Web server).
- 3: the number of generated processes is out of thresholds, and we don't have any tool (script) to check whether the service is operative (this could be the case of processes such as syslogd, crond and xinetd).
- 4: service is not responding to requests within the defined time period.

We group the above five situations in three more general cases:

- OK (status 0 and 1).
- WARN (status 2).
- DOWN (status 3 and 4).



axigen

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

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

Logos for Linux, Solaris, FreeBSD, and NetBSD are shown.

FIRST SUPPORT
(FREE IMMEDIATE RESPONSE TIME SUPPORT)

AVAILABLE 24x7 FOR ALL OUR USERS!
ALSO AVAILABLE DURING EVALUATION!

AXIGEN MAIL SERVER

www.axigen.com

Register now in less than 1 minute
for the FREE OFFICE Edition

Listing 3.

Sample Output of the smart -d Command

```
[sysman@server ~]# ./smart -d
```

SERVICE	PID	PROCS	STATUS	PROBLEM
CRON	451	1	[OK]	
DISK	?	0	[OK]	No start command.
DHCP	444	1	[OK]	
DNS	442	1	[OK]	
HTTP	625	53	[WARN]	Too many processes (>30).
LPD	474	1	[OK]	
MRTG	27017	1	[OK]	
MYSQL	627	1	[OK]	
NAGIOS	640	1	[OK]	
NMB	633	1	[OK]	
NTP	?	1	[OK]	
POSTFIX	619	0	[DOWN]	No response from service. [Starting...]
->POSTFIX	23945	1	[OK]	
POSTGRES	560	3	[OK]	
SLAP	643	1	[OK]	
SMB	631	6	[OK]	
SNMP	635	1	[OK]	
SNMPTRAP	637	1	[OK]	
SSH	654	3	[OK]	
SYSLOG	402	1	[OK]	
XINET	462	1	[OK]	

Listing 4.

A Sample of the nag and Shell Scripts

```
[root@server /]# ls /home/sysman/scripts/
```

disk.nag	http-forb.nag	nfs.nag	pop3.nag	smtp.nag
disk.sh	http.nag	nfs.sh	printer.nag	snmp.nag
dns.nag	http.sh	nmb.sh	proxy.nag	ssh2.nag
dns.sh	imap.exp	ntp.sh	slap.nag	ssh.nag
ftp.exp	imap.nag	pgsql2.nag	slap.sh	ssh.sh
ftp.nag	mysql.nag	pgsql.nag	smb.nag	
http-auth.nag	mysql.sh	pgsql.sh	smb.sh	
http.exp	nagios.nag	pop3.exp	smtp.exp	

Listing 5.

nag scripts are handled by plugins.

```
[root@server /]# ls /home/sysman/plugins/
```

check_disk	check_http	check_pgsql	check_snmp	check_udp
check_dns	check_imap	check_pop	check_ssh	
check_ftp	check_nagios	check_smtp	check_tcp	

Decision

When executing the program with no parameters, it simply will determine the status of services defined in the configuration file and will display the results. If we want the program to work in an active way, we need to use some of the following parameters:

- **-w**: restart services in WARN status and send a notification (e-mail) for each one of them.
- **-d**: restart services in DOWN status and send a notification for each one of them.
- **-wd**: restart services in WARN and DOWN status and send a notification for each one of them.
- **--all**: restart all services independently of their status and send a notification for each service with WARN or DOWN status.
- **--reboot**: restart the whole system independently of service's status and send a general notification.

Once the service status has been determined, and according to the parameter specified in the execution, the action carried out for each service will consist of that shown in Table 1.

Furthermore, independently of the service's status, with the parameters **--all** and **--reboot**, a notification via e-mail is sent to the administrator about the performed action.

Listing 3 shows a sample of SMART in action, executed from a console with parameter **-d** (recovery of services in DOWN status).

Check Scripts

There are some optional executables files, the check scripts, responsible for checking whether the monitored services really are operative and responding to petitions. These files are written in Shell (.sh extension) and Expect (.exp extension). Expect is a tool that requires Tcl and allows for automation of interactive applications that use textual representation.

These scripts could be written in any programming language, because only the exit status is taken into account. If it's not equal to 0, we suppose that there has been no answer or that the answer given by the service has not been the expected one. This means that a check script not only can monitor services, but it also can achieve any check that returns a Boolean value, for example, to check whether the size of a directory exceeds a certain value, whether the amount of logged users is greater than a desired number, whether a kernel module is loaded and so on (Listing 4).

Files with the .nag extension are also Shell scripts, but unlike the former ones, they call an external program (plugin) passing to it the parameters received from check-service, following the order and format that the plugin expects. This

Table 1. Service Actions

Status	Parameters	Action
OK	--all	Restart the service
WARN	-w, -wd, --all	Restart the service
		Send a notification relating to service
	-d	Send a notification relating to service
DOWN	-d, -wd, --all	Restart the service
		Send a notification relating to service

checks the service and returns the information gathered to the check script, which will interpret and convert it into the exit status that check-service is waiting for (Listing 5).

Plugins are programmed in C, Perl and Shell and belong to Nagios. Their sources can be downloaded independently of the Nagios distribution, and some of them require the additional installation of certain programs and libraries.

Installation, Configuration and Usage

Software requirements include the following:

- **sudo**: allows a user to execute a command as another user. This will be necessary if you are planning to allow a nonroot user to execute SMART.
- **awk**: a pattern scanning and processing language. SMART uses it and expects to find it at /bin/awk. If that's not your case, edit the check-service and smart files of the SMART distribution and modify the line where AWK="/bin/awk" is specified.
- **Nagios plugins**: sources can be downloaded independently of the Nagios distribution, and some of them require the additional installation of certain programs and libraries. You can use the plugins distributed with SMART or download the newest ones.
- **Some shell scripts** (in the scripts directory of SMART) may require some specific commands to check some services, such as dig for dns, wget for Web services, nmblookup for nmbd (Samba), ntpq for NTP, ldapsearch for OpenLDAP and so on. The paths of these commands are defined in a variable at the beginning of each script, so you can change their location, use any other command that might work better for your system or even rewrite the whole script at your convenience.

With sudo you can permit another user to run SMART. If you're not interested in creating such a user, you can omit steps 1, 2 and 3 below.

1. Create user sysman and group sysman.
2. Create the SMART directory. It's a good idea to install it at sysman home and to set the appropriate owner and permissions:

```
mkdir /home/sysman
chown root:sysman /home/sysman
chmod 750 /home/sysman
```

3. Edit the sudo configuration file /etc/sudoers, and add the following lines:

```
...
sysman hostname=(root) NOPASSWD: /home/sysman/check-service
sysman hostname=(root) NOPASSWD: /sbin/reboot
```

4. Download the SMART software.

5. Untar and unzip the distribution:

```
tar -zxf smart-X.Y.tar.gz
```

6. Go to the distribution directory and copy the files to the destination directory. If you choose a destination different from /home/sysman, you will have to edit the smart file and modify the line where dir="/home/sysman" is specified:

```
cd smart-X.
cp check-service /home/sysman/
cp smart /home/sysman/
cp host.conf.dist /home/sysman/host.conf
cp services.conf.dist /home/sysman/services.conf
cp -r scripts /home/sysman/
cp -r plugins /home/sysman/
```

7. Go to the destination directory, and check/set file permissions and owners:

```
cd /home/sysman
chown -R root:root check-service scripts plugins host.conf services.conf
chown root:sysman smart
chmod -R 700 check-service scripts plugins
chmod 750 smart
chmod 644 host.conf services.conf
```

Configuration is as follows. First, edit the SMART host configuration file host.conf, and modify it according to your preferences (hostname, mail addresses, commands paths and so on). Then, edit the SMART services configuration file services.conf, and uncomment/modify/add any service/dæmon you want to check. Every line describes one service, with the following semicolon-separated parameters:

- **NAME** (non-empty string): descriptive service name (for example, IMAP).
- **process_name[:port]** (non-empty string[:integer]): parent process name and its operational port (for example, couriertcpd:143).
- **process_param** (string): parameters of running process. Some services run with the same process name, so parameters are useful to distinguish them. For example, the parent process of Courier IMAP and POP3 is couriertcpd, but one is executed with the parameter pop3d and the other one with imapd.
- **max_procs** (non-empty integer): the highest number of running processes allowed (for example, 10). Leave it at 0 if what you're monitoring runs no processes (for example, disk space).

Do you take

"the computer doesn't do that"

as a personal challenge?

So do we.

LINUX
JOURNAL™

Since 1994: The Original Monthly Magazine of the Linux Community

Subscribe today at www.linuxjournal.com

- `min_procs` (non-empty integer): the lowest number of running processes allowed (for example, 1). Leave it at 0 if what you're monitoring runs no processes (for example, disk space).
- `start_command` (string): the command to start the service or script to be executed when the service is down (for example, `/courier/libexec/imapd.rc`).
- `pid_file` (string): pid file path (for example, `/var/run/imapd.pid`).
- `sock_file` (string): socket file path.
- `start_mode` (0/1): the service can be started/stopped by adding start/stop to the start command (1), or it may not be necessary (0).
- `check_script` (string): the name of the script used to check the service. This script has to be in the scripts directory (for example, `imap.nag`).

Leave the parameters empty if they are not applicable, except `NAME`, `process_name`, `max_procs`, `min_procs` and `start_mode`, which can't be empty.

Now, you should be able to run SMART as user `root` or `sysman`:

```
/home/sysman/smart
```

Try using `-h` to get more information about available parameters. Running SMART through `crond` might be a good thing. You can run it as frequently as you want, but doing it every five minutes seems to be reasonable enough.

Conclusion

SMART is an easy-to-install application (simply copying the program), is much simpler to configure than Nagios (adding a new element to monitor involves adding only one line in the configuration file), and SMART is flexible, allowing you to monitor any service or aspect of the system, and it is very effective.

Our experience in a production environment with thousands of users tells us that it's inevitable that we will reach some peak periods in which the amount of requests received by a service goes beyond the capabilities of the system, and response time grows in a dramatic manner. The fact that the system detects this situation, before its own administrator, and solves it in five minutes, is a great problem solver and provides a perception of better service to users.

After two years of running SMART on about 15 servers, we can say that its main contribution has been our peace of mind. It's wonderful having a colleague who is checking that everything works correctly 24/7 and who informs you about troubles after they already have been solved (especially during the weekends).

Acknowledgements

SMART was created, developed, tested and enjoyed in the IT Department of the Universitat Internacional de Catalunya. Vicente Sangrador and Jordi Xavier Prat have collaborated on this project and encouraged me to write this article. ■

Resources for this article: www.linuxjournal.com/article/9268.

Albert Martorell is a Telecommunications Engineer and has been working as a network and "penguins" administrator in the IT Department of the Universitat Internacional de Catalunya since 1998.

Advertiser Index

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

Advertiser	Page #	Advertiser	Page #
AMAX www.amax.com	23	MBX www.mbx.com	21
APPRO HPC SOLUTIONS appro.com	C2	MICROWAY, INC. www.microway.com	C4, 91
ASA COMPUTERS www.asacomputers.com	27, 85	MIKRO TIK www.mikrotik.com	7
AVOCENT CORPORATION www.avocent.com/special	1	MONARCH COMPUTERS www.monarchcomputer.com	10, 11
CARI.NET www.cari.net	71	OFFASSIST www.offassist.com	42
CIARA TECHNOLOGY www.ciara-tech.com	8, 9	OPEN SOURCE SYSTEMS, INC. www.opensourcestorage.com	13
CORAID, INC. www.coraid.com	37	OPERA SOFTWARE www.opera.com	5
COYOTE POINT www.coyotepoint.com	3	PENGUIN COMPUTING www.penguincomputing.com	47
EMAC, INC. www.emacinc.com	52	POGO LINUX www.pogolinux.com	53
EMPERORLINUX www.emperorlinux.com	15	POLYWELL COMPUTERS, INC. www.polywell.com	59
FAIRCOM CORPORATION www.faircom.com	63	THE PORTLAND GROUP www.pggroup.com	31
GECAD TECHNOLOGIES/AXIGEN www.axigen.com	77	RACKSPACE MANAGED HOSTING www.rackspace.com	C3
GENSTOR SYSTEMS, INC. www.genstor.com	69	R CUBED TECHNOLOGIES www.rcubedtech.com	73
HURRICANE ELECTRIC www.he.net	29	SERVERS DIRECT www.serversdirect.com	43
INTEL www.intel.com/software/products	25	SILICON MECHANICS www.siliconmechanics.com	45, 57
IRON SYSTEMS www.ironsystems.com	87	SUPERMICRO www.supermicro.com	35
ISPCON www.ispcon.com	34	TECHNOLOGIC SYSTEMS www.embeddedx86.com	65
JTL NETWORKS www.jtlinet.com	6	THINKMATE www.thinkmate.com	17
LANTRONIX www.lantronix.com	41	USENIX LISA www.usenix.org/events/lisa06	89
LAYER 42 NETWORKS www.layer42.net	90	@XI COMPUTER CORPORATION www.xicomputer.com	19
LINUX JOURNAL www.linuxjournal.com	53	ZT GROUP INTERNATIONAL www.ztgroup.com	33
LPI www.lpi.org	75		

A Basic Text-Based Recording Studio

Forget the huge expensive mixers and create a recording studio in a Linux box.

MATTHEW GEDDES

Whether you're into Metal, Jazz, Noise, Baroque or something in between, it is becoming more and more popular for artists to take on not only the roles of composer and performer, but also the roles of audio engineer, producer and even distributor of their own work.

The capability and quality of Linux audio applications are very good and constantly improving. Support for high-end and low-end audio cards is also getting better all the time. Whether it becomes the dominant platform in the field is largely irrelevant—those of us who find the flexibility of Linux and open-source tools to be valuable now have a platform suitable for creating high-quality audio tracks.

This article outlines a simple method, which may be built upon, for recording layered, multitrack recordings. In keeping with the Linux tradition, in this article, we discuss a number of small, command-line tools that perform very specific tasks very well. We then combine the power of each of these tools into a digital audio workstation.

As you will see, using these tools in such a way, it is possible to overcome the (rare) shortcomings in some of these tools.

The tools we cover here are Ecasound and JACK. The Hydrogen drum machine is mentioned briefly too. We use no ALSA- or OSS-specific features directly, and either will do fine. In fact, for those who have lost their way and have strayed from the path to enlightenment (kidding), these tools and techniques also work under CoreAudio on Mac OS X.

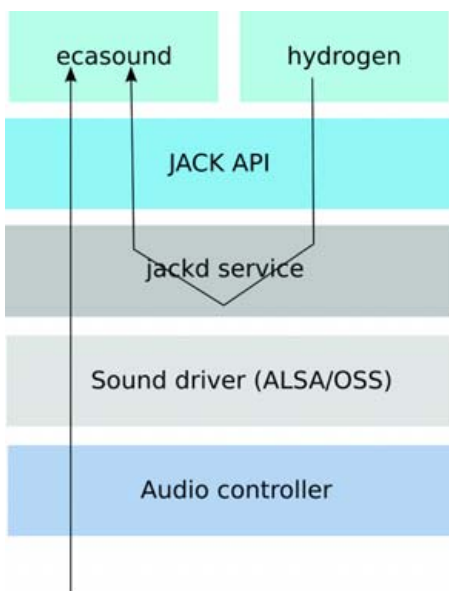


Figure 1. Audio Data Flow

Figure 1 shows how data flows between each of these components at a high level.

Equipment List

For the examples outlined in this article, any sound card will do. I even have performed some relatively acceptable recordings using the onboard Intel i8x0 sound device in one of my Linux laptops. However, the difference between lower-end audio controllers and the mid- to

high-end ones is quite noticeable.

We also require a Linux distribution. If you have trouble getting JACK and Ecasound for your distribution, try the AGNULA live distribution. Most distributions come with the relevant packages these days anyway.

A mixer is desirable. Using a small (read: cheap) mixer may give you more flexibility and a chance at a better sound. You also may find that a direct injection box or a microphone preamp is adequate.

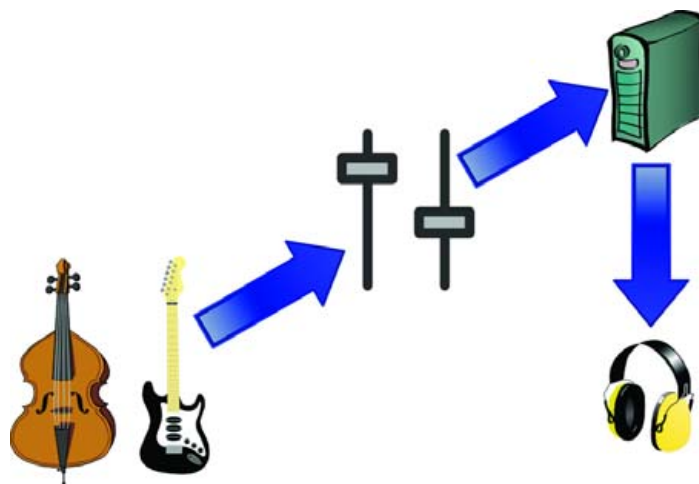


Figure 2. Equipment Signal Flow

Note that Figure 2 suggests plugging the headphones in to your Linux box. Most mixers allow the sound card to be plugged in to a signal return port and allow the headphones to receive the audio signal either before or after the signal is sent to the Linux box. This is acceptable too.

Assumed Knowledge

In this article, we don't assume much apart from the following:

- A Linux box with a configured and tested audio controller.
- The ability to source and install necessary packages and their dependencies.
- A familiarity with your choice of noise-making device (for example, guitar, cello, cat and so on).

Starting jackd

JACK, which stands for the JACK Audio Connection Kit, is an API and a service that provides audio connectivity between applications on

many POSIX-compliant systems. JACK has been designed with low-latency communication in mind.

Many of the examples in this article may work equally well without JACK. I personally have had fewer audio dropouts on systems employing JACK running with real-time priority than without, and it is quite useful for interconnecting audio applications, such as the Ecasound and Hydrogen example discussed later.

For applications to make use of JACK, they must be linked against the JACK API libraries, and the JACK service, called `jackd`, must be started. Distributions shipped with JACK often already have most applications linked against the JACK API. If not, consult the build or compile instructions for your given application.

To start the JACK service, execute a command similar to the following:

```
jackd -R -d alsa
```

The `-R` option instructs JACK to attempt to attain real-time privilege, and `-d alsa` instructs JACK to use the ALSA sound system. For users still using the OSS sound system, `-d oss` should suffice, and `-d coreaudio` should get Mac OS X users off to a start.

Each driver supports a series of driver-specific options. These may be viewed by specifying `--help` after `-d alsa`.

Testing Audio Signal and Setting Levels

Before leaping in too far and beginning to record audio, I strongly recommend spending some time getting the various settings and levels right. The good news is that this involves plugging the instrument of your choice in to the mixer or sitting it in front of a microphone and playing.

Begin by getting the average signal coming into the mixer at around the 0VU mark and try to avoid sending the signal into the red too often. Once the mixer levels are generally okay, connect it to your PC and check that the input level and output level are fine:

```
ecasound -i jack_auto -o null -ev
```

The `-i jack_auto` command-line option tells Ecasound to get its input from JACK. Because we're not running any other JACK-aware applications at the moment, JACK takes this input from the sound device. The `-o null` tells Ecasound to send output to the great bit bucket in the sky.

The `-ev` option tells Ecasound to keep track of amplitude statistics, and the `-c` option starts Ecasound in interactive mode. With a little luck, you should see a few informational messages and no errors or warnings.

Any percussive sounds (such as palm-muting on the guitar) are likely to cause a spike in your audio track. While checking the signal levels, use any of these techniques you intend to record later—it'll save a nasty surprise in the moment of creative genius. To stop, press `Ctrl-C`. You should be presented with output similar to the following:

```
....
(audiofx) Peak amplitude, period: pos=0.30495 neg=0.26996.
(audiofx) Peak amplitude, all : pos=0.30495 neg=0.26996.
(audiofx) Clipped samples, period: pos=0 neg=0.
(audiofx) Clipped samples, all : pos=0 neg=0.
(audiofx) Max gain without clipping, all: 3.27926.
(audiofx) -- End of statistics -----
```

First, check that you have no clipped samples (positive or negative).

Second, check the maximum gain figure. This gives the percentage that this sample can be amplified (theoretically) before clipping starts to occur. Depending on your hardware, you may never get within a few percentages before you hear audible distortion, so it pays to leave yourself a little room until you're familiar with your hardware. Listen as you test.

Once you have made mixer adjustments, try the previous few steps again.

Once you're happy with the input levels, set the output level to a comfortable level for you to monitor using your headphones.

Recording a First Track or Live Stereo Performance

Ecasound is a command-line tool capable of multitrack recording and more. The basic concept key to using Ecasound is chains. For our purposes, you can consider chains to be similar in function to a patch lead in a patch bay. A signal enters one end of the chain from a sound source and exits the chain into another component. A patch lead has exactly one input source and one output destination, and the same can be said about Ecasound's chain concept.

Sources and destinations for chains in Ecasound are usually audio files or audio controllers. It is quite normal to have a complex set of chains. The first track we will record will see Ecasound take audio from the running JACK instance and write the data back to JACK, as well as keep a copy in a PCM audio file. The two chains we need to perform these tasks are shown in Table 1.

Table 1. Chains for Our First Track

Chain	Input Source	Output Destination
1	JACK	JACK
2	JACK	track1.wav

This equates to the following Ecasound command:

```
ecasound -c -b:64 \  
-a:1,2 -i jack_auto \  
-a:1 -o jack_auto \  
-a:2 -o track1.wav
```

Once Ecasound has initialised, it prompts you for instructions. Use the `t` command to start recording/playing and `s` to stop. If you make a mistake, you can issue a stop (`s`), the `setpos 0` command, and `t` to start again. The `q` command quits when you're done. There's no need to issue any kind of command to save the result—that happens as you record.

The above command can be broken down into the following functions:

- `-c`: don't start processing immediately, instead enter interactive mode.
- `-b:64`: set the number of samples buffered to the smallest possible, reducing latency.
- `-a:1,2 -i jack_auto`: create two chains (1 and 2) and set their input to come from JACK.

- -a:1 -o jack_auto: set the output of chain 1 to JACK.
- -a:2 -o track1.wav: set the output of chain 2 to track1.wav.

The overall result of this particular example is that chain 2 records anything coming in through JACK (and therefore probably the sound card) to track1.wav. Chain 1 allows you to hear the audio signal as it's being recorded.

Overdubbing of Subsequent Tracks

Unless you're recording a live stereo track, you're likely to want to overdub other tracks. It is possible to use Ecasound to listen to tracks you've already recorded while recording (and listening to) a new track.

To listen to an already-recorded track while recording a second track, create three Ecasound chains (Table 2).

Table 2. Chains for Listening to an Already-Recorded Track While Recording a Second Track

Chain	Input Source	Output Destination
1	track1.wav	JACK
2	JACK	JACK
3	JACK	track2.wav

Creating a chain setup like this causes the contents of track1.wav to be sent to JACK to be played, and any input from JACK is sent back to JACK and saved to a file called track2.wav. track2.wav will contain only the new track—not the new track mixed with the old track. We'll mix them later.

Converting this chain setup to an actual Ecasound command is straightforward:

```
ecasound -c -b:64 \
-a:1 -i track1.wav \
-a:2,3 -i jack_auto \
-a:1,2 -o jack_auto \
-a:3 -o track2.wav
```

Recording further tracks is a similar process. We create a chain for each of our already-recorded tracks and set their output to JACK. We also set up two chains to take input from JACK and send it to a file and back to JACK, so we can hear it. The chain setup in Table 3 would suffice.

Table 3. Chains for Recording More Tracks

Chain	Input Source	Output Destination
1	track1.wav	JACK
2	track2.wav	JACK
3	JACK	JACK
4	JACK	track3.wav

This chain setup translates into the following Ecasound command:

```
ecasound -c -b:64 \
-a:1 -i track1.wav \
-a:2 -i track2.wav \
-a:3,4 -i jack_auto \
-a:1,2,3 -o jack_auto \
-a:3 -o track3.wav
```

Recording a Software-Based Sound Source

I'm not a drummer, but some of the things I record need drums. Although the Hydrogen drum machine is probably the best that I have seen on Linux, it hasn't yet attained the magical 1.0 version number and isn't yet perfect. One feature that's broken in the snapshot I'm running is the ability to export to a PCM audio .wav file. As luck, or rather good design, would have it, Hydrogen can use JACK to output digital audio.

To use JACK and Ecasound to record the output of an audio application such as Hydrogen, we can perform the following steps:

1. Configure Hydrogen to use JACK for its output.
2. Configure Hydrogen to play in song mode, as opposed to pattern mode.
3. Execute the the following command:

```
ecasound -a:1 -i jack_auto -o drum_track.wav -G jack,ecasound,recv
```

4. Click the play button in Hydrogen.

The above command configures a single chain within Ecasound that draws input from JACK and sends output to a file called drum_track.wav. The -G jack,ecasound,recv instructs Ecasound to listen to JACK for a start command, which is sent when we click the play button in Hydrogen, as a JACK client called ecasound.

It can take a second or so for Ecasound to start and initialise after receiving the start command, so I like to have a pattern of silence at the start of the Hydrogen track.

Unless you have impeccable timing, you would most likely record any software sources first. It is harder to synchronise a software source, such as a drum machine, with an existing human-recorded track than it is to record the human tracks around the machine-created tracks. This includes any MIDI tracks you intend to use.

Mixing All Tracks to a Single Stereo Master

At this point, we have a series of .wav files that correspond to each of the audio tracks we have recorded. Should we need to, we could use Ecasound, SoX or even Audacity to add effects or make

Table 4. Turning Multiple Tracks into a Single Stereo Master

Chain	Input Source	Output Destination
1	track1.wav	all_tracks.wav
2	track2.wav	all_tracks.wav
3	track3.wav	all_tracks.wav

You also can attach a series of effects, including reverb, compression and amplification to each chain before it is written to the output destination.

minor corrections or alterations to any of the tracks. Once we're happy with the individual tracks, we can mix a single master track.

The process of turning our multiple tracks into a single stereo master track is straightforward. We create a chain for each track and set the output to be a .wav file.

Ecasound provides a means to make this particular case easier. The all pseudo-chain name can be used to redirect the output of all of our tracks to a single place, namely a file called all_tracks.wav:

```
ecasound -a:1 -i track1.wav \  
-a:2 -i track2.wav \  
-a:3 -i track3.wav \  
-a:all -o all_tracks.wav
```

To listen to the result, enter:

```
ecasound -a:1 -i all_tracks.wav -o jack_auto
```

It is also possible to write the master directly to the sound card without writing to a file first:

```
ecasound -a:1 -i track1.wav \  
-a:2 -i track2.wav \  
-a:3 -i track3.wav \  
-a:all -o jack_auto
```

You also can attach a series of effects, including reverb, compression and amplification to each chain before it is written to the output destination. It is even possible to add delay and alter the panning of a particular track or even perform noise reduction; however, such topics are beyond the scope of this article.

Summary

As we have demonstrated, it is possible to create a simple multitrack recording using a handful of Linux audio tools. Once we started jackd, it was a simple process of telling Ecasound where to receive input from and where to send output to as we recorded our initial track and overdubbed a series of subsequent tracks.

Each of these tracks has been stored in its own individual .wav file. This allows us to use any other soundfile editor to make manual modifications to the track before mixing a final track, which can then also be tweaked. Common applications for processing audio files include Ecasound, SoX and Audacity.

We have really just scratched the surface of this particular aspect of a large field. With luck, it will form a solid foundation on which you can build your creative genius! ■

Resources for this article: www.linuxjournal.com/article/9269.

Matthew Geddes' hobbies are music and Linux. Luckily for him, and those around him, they also happen to be his career. When he's not playing his own stuff, he's listening to everything from Bach and Son House to Rachel Singleton and A norexia Nervosa. He can be reached at lj@musicalcarrion.com or through www.musicalcarrion.com.



ASA COMPUTERS

Want your business to be more productive?

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

Hardware Systems For The Open Source Community—Since 1989

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

Dempsey/Woodcrest Server

\$2024

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



Dual Xeon Server starts at

\$4015



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

Dual Xeon 800FSB Storage starts at

\$6,644


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



Your Custom Appliance Solution

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

"Your Logo Here"



ASA Collocation

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

ASA Collocation Special

First month of collocation free.*

Storage Solutions

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

Building and Integrating a Small Office Intranet

This “how we did it” story includes valuable tips for building an intranet that integrates enterprise services in a user-friendly way.

DAVE JONES

Intranets have been around for a long time. They were one of the first alternate uses for World Wide Web technology back in the early 1990s. The idea of bringing a little bit of the Web experience in-house was very attractive, but integration with existing systems was difficult. Thus, a lot of intranets were nothing more than glorified bulletin boards with some user-publishing features thrown in. The landscape is different now, with open-source software ready to take most of the cost and some of the complexity away from a good intranet setup. The so-called LAMP stack provides the perfect neutral platform for integrating many different pieces of software into a single point of interaction for users. That’s what we have tried to do at our company.

Our intranet started off in 1999 as a Web-based bulletin board and company calendar on a Red Hat 6.0 server running Apache. It was a static HTML site that was designed and kept current by our marketing manager. After she left the company in 2002, we needed to make the intranet more dynamic so that it didn’t depend on one person to keep it up to date. As is usually the case, we added more and more features over the years and now have a very useful, user-friendly intranet site without a lot of unnecessary or static content that needs to be maintained. In this article, I use our intranet as an example of how to solve four of the more common integration tasks that small business admins may run into when setting up a LAMP-based intranet.

Technical Overview

Our intranet currently serves about 70 employees and runs on an IBM x335 server running Fedora Core 4. We use a normal LAMP stack (Linux, Apache 1.3x, MySQL and Perl) with mod_perl to improve performance. Apache currently shares the server with our e-mail scanner, internal DNS, Jabber, Samba and some other services. It’s nice having all of this running on a single Linux server, because it reduces the need for NFS mounts and cuts down on network traffic. Some sites will be too large for this approach, but nothing in our design would preclude

it from working in a multiserver setup. All of our users run Windows XP and authenticate through Active Directory. We use GroupWise as our e-mail software running on a NetWare 6 server, and all of its information is handled by Novell’s eDirectory. We also have a time and billing system that runs on a Windows NT 4.0 server and stores its data in a Microsoft SQL Server database. You can see a layout of how everything links together in Figure 1.

Server-Side Credentialing

We decided early on that our users shouldn’t have to authenticate in any way to our intranet. The site should automatically “know” who they are based on their IP address and information gleaned from the network about who is currently logged in from that address. We call this technique server-side credentialing (SSC). We accomplished this originally by using a piece of custom-written client-side software that was contacted by a CGI script any time the server needed to check a user’s identity. This works, but it places too much trust on the client side. A sniffer and a Perl script could fake a user’s identity nicely from any client computer. We now use Samba and winbindd for this task.

Because our intranet server resides on a trusted internal network, it is privy to the current state of affairs on the network, including who is logged in from where. Every computer in the office maps a drive letter to the Samba server during login, so any time the server needs the current user’s identity, it simply looks up his or her IP address in the Samba connection list. The mapped drive is just a dummy drive explicitly for the SSC mechanism. I think this is an important feature, because it lowers the complexity of the site from a programming standpoint and allows users to browse freely without having to worry about registering or logging in. Users have enough user names and passwords to keep track of already without us adding to their burden.

The way you set up SSC depends on how your users authenticate on your network. We use Active Directory, so that is what I demonstrate here. Active Directory is annoying (surprise, surprise), because it doesn’t store connection status information in its directory. You must use traditional RPC calls with Samba’s net command to get reliable results. Our SSC script is called smbconn.sh, and it looks like this:

```
#!/bin/sh

net status sessions parseable \
| grep -i "\\\\$1\\\\" \
| sed 's/^\.*\\(.*)\\.*\\.*\\.*$/\1/g' \
| sed 's/DOMAIN+//g' | tr -d ' '
```

Pretty simple, eh? Just remember to change DOMAIN to whatever your Active Directory’s domain name is. This script returns the name of the user object that is logged in to Samba from the IP address we pass to it on the command line. The name it returns corresponds to Active




Figure 1. How Our Enterprise Services Are Connected

Directory's sAMAccountName property. Armed with this information, we now can run an LDAP lookup to get the user's full name or any other data we might need. The script we use to do this is found in Listing 1. It will take the user's sAMAccountName as its first argument and an optional attribute whose value you want returned as the second argument. If you don't provide the optional attribute, the script returns the user's full name. You *could* do all of this in a custom mod_perl handler so that its information always would be available, but this seems like overkill for most sites. Our site has only a handful of restricted sections where this information comes into play, so we just let each CGI script run it as needed. Here is a typical SSC call from one of our CGI scripts:

```
##: Get this connection's user credentials
my $remoteip=$ENV{'REMOTE_ADDR'};

open(SMBCONN,"smbconn.sh $remoteip |");
my $cfn=<SMBCONN>;
$cfn=~s/\s+//g;    ##: Strip whitespace
close(SMBCONN);

open(GETEMPINFO,"getempinfo.pl $cfn |");
my $username=<GETEMPINFO>;
close(GETEMPINFO);
if($username eq "") {
    $username="Guest";
}
}
```

This section of code leaves us with the user's sAMAccountName in the \$cfn variable and the user's full name in the \$username variable. If the \$username variable contains Guest, either the lookup failed or the computer accessing this CGI script doesn't have a logged-in user operating it. We now can use this critical information to decide whether the user has access to the information this CGI script is meant to provide. We also can use this information to return a page customized for this particular user. I demonstrate this with a section of code from the index.cgi file that serves up our home page:

```
##: My Intranet section
my $mint="";
if(($username eq "Guest") || ($username eq "")) {
    open(EMPSNAP,"./random-employee.pl 2>&1 |");
    my @snap=<EMPSNAP>;
    close(EMPSNAP);
    $mint.=join("\n",@snap);
    chop($mint);
} else {
    $mint.=&get_emp_card($cfn);
    $mint.="<b>E-Mail Controls:</b><br>\n";
    $mint.="<a href='selfserv.cgi'>My Mail</a>\n";
    ...
}
...
print STDOUT $mint;
```

You can see here that we check to see if the person viewing the home page is actually a credentialed user. If he or she is not, we serve up a random employee's picture and profile in this section of the home page. If the person is a credentialed user, we grab the

appropriate personal information from the LDAP directory and proceed to assemble a My Intranet area in this section of the home page where the user can edit his or her employee profile, control mail preferences and so forth. The get_emp_card(\$cfn) routine simply looks up the user's current info in Active Directory and returns a nicely formatted HTML section to display it (Figure 2).

Active Directory Integration

Another valuable addition to our intranet was integrating it with our Active Directory user database via LDAP. We use this to provide a company directory that lists all of our employees. The directory is built in real time whenever it is accessed, and that is a major time-saver for administrators. Whenever new users are added using the normal Active Directory tools, they instantly show up in the intranet directory. We also allow our users to edit their own personal information, and those edits are put into the Active Directory by the CGI script. The process is relatively straightforward, although there are some things to take into consideration. Let me walk you through the process of how we set this up.

The first thing we do is create a proxy user called proxyuser in Active Directory. This is the user name our scripts use to authenticate with LDAP. The proxy user is granted rights to read and write information on user objects within the ou=Domain Users container. That's all that needs to be done within Active Directory. We use Perl for our CGI, so that means using Net::LDAP. Here is how we connect to Active Directory from within a CGI script:

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

Listing 1.

getempinfo.pl

```
#!/usr/bin/perl -w
use Net::LDAP;
use strict;

my $cn=$ARGV[0] || "none";
my $attr=$ARGV[1] || "none";

##: If nothing was given on command line then return
if($cn eq "none") {
    print STDERR "ERROR: No LDAP cn given\n";
    exit(1);
}

##: Bind anonymously to the ldap database
my $ldap=Net::LDAP->new('directory.domain.com',timeout=>5)
    or die "Couldn't connect to directory server.\n";
my $mesg=$ldap->bind('proxyuser@domain.com',password=>'proxyuser')
    or die "Couldn't connect to directory server.\n";

##: Query LDAP to get a list of employees
if($attr ne "none") {
    $mesg=$ldap->search( base=> "ou=Domain Users,dc=domain,dc=com",
        filter=> "(sAMAccountName=$cn)",
        attrs=> ['givenName','sn',$attr] );
} else {
    $mesg=$ldap->search( base=> "ou=Domain Users,dc=domain,dc=com",
        filter=> "(sAMAccountName=$cn)",
        attrs=> ['givenName','sn'] );
}

my $count=$mesg->count();
($count==1) or die "Error: LDAP enumeration error.";

my $entry=$mesg->entry();
my $value;
my @values;
if($attr ne "none") {
    $value="";
    @values=$entry->get_value("$attr");
    my $i=1;
    for(@values) {
        if($i>1) {
            $value.="/$_";
        } else {
            $value=$_;
        }
        $i++;
    }
} else {
    $value=($entry->get_value('givenName')." ";
    $value.=$entry->get_value('sn'));
}

##: See if that attribute was defined for the given cn
if(!(defined($value))) {
    print STDERR "ERROR: That attribute was not defined.\n";
    exit(1);
}

$mesg=$ldap->unbind;
print("$value\n");
```

Figure 2. A Sample User Profile on the Intranet

```
##: Active Directory connection
use Net::LDAP;

my $ldap=Net::LDAP->new('adserver.domain.com');
my $mesg=$ldap->bind('proxyuser@domain.com',
    password=>'proxyuser' );
```

Notice the syntax that Active Directory requires for the user name field. It's one of the unique requirements of Active Directory's LDAP interface. Now that we are connected to the directory, we do a query to find all the user objects in the ou=Domain Users container:

```
##: Query LDAP to get a list of employees
my $basedn="ou=Domain Users,dc=domain,dc=com";
my $filter="(objectClass=user)";
$mesg=$ldap->search(
    base=> $basedn,
    filter=> $filter,
    attrs=> ['givenName','sn','mail',
        'telephoneNumber','streetAddress',
        'l','st','department','postalCode',
        'employeeNumber','homePhone',
        'title','sAMAccountName' ]
);
```

Sponsored by

THE USENIX SIG FOR
USENIX & **[sage]**
SYSADMINS

6 DAYS OF TRAINING
BY INDUSTRY EXPERTS, INCLUDING:

- Gerald Carter on Ethereal and the Art of Debugging Networks
- Richard Bejtlich on TCP/IP Weapons
- Aileen Frisch on Administering Linux in Production Environments
- Chip Salzenberg on Higher-Order Perl
- And 55 other tutorials

LISA '06

20TH LARGE INSTALLATION
SYSTEM ADMINISTRATION CONFERENCE

A **Blueprint** for Real World
System Administration

DECEMBER 3-8, 2006 | WASHINGTON, D.C.

3-DAY TECHNICAL PROGRAM

Keynote: Cory Doctorow, science fiction writer, co-editor of Boing Boing, and former Director of European Affairs for the EFF, on Hollywood's Secret War on Your NOC

20+ Invited Talks, including:

- Simple Nomad, Vernier Networks, Inc., "Corporate Security: A Hacker Perspective"
- DJ Byrne, Jet Propulsion Laboratory, "Open Source Software and Its Role in Space Exploration"
- Mazda Marvasti, Integrien, "Everything You Know About Monitoring Is Wrong"

Refereed Papers, Hit the Ground Running Track, Guru Is In Sessions, Vendor Exhibition, Workshops, BoFs, WiPs, and more!

Register by November 10 and save! www.usenix.org/lisa06/lja

LISA '06 offers the most in-depth, real-world system administration training available!

This returns all of the user objects in that container, along with all of the pertinent attributes you would expect to find in a company directory. We now can refine our search filter to limit our search to only those users whose last name starts with a letter passed to the CGI script in its URL. This allows us to follow an address-book format, so we don't have to display all 70 users at once. We fall back to the letter a if no letter was asked for in the URL:

```
##: Get letter requested in the URL
my $letter;
$letter=param('letter') || "a";
...
my $filter="(&(objectClass=user) (sn=$letter*))";
```

If you aren't familiar with the syntax used by LDAP search filters, I suggest you look over RFC-2254. At this point, we can iterate over our query results and prettify them as needed. Because we also looked up this user's SSC information, we can check each employee's sAMAccountName as we go through the loop. When we find the employee that corresponds to the person SSC says is viewing the page, we add a link by the employee's name that allows him or her to go to an area to edit the directory information. It looks like this:

```
##: Display the directory
foreach my $entry ($msg->sorted('sn')) {

    my $san=$entry->get_value('sAMAccountName');
    $empdir.="<div class='empcard'>";
    if(lc($cn) eq lc($san)) {
        ##: This is our man. Add a button.
        $empdir.="<a href='empedit.cgi'>Edit</a>";
    }
    $empdir.="<span id='name'>";
    $empdir.=$entry->get_value('givenName')." ";
    $empdir.=$entry->get_value('sn');
    $empdir.="</span><br>";
    $empdir.="<span id='title'>";
    $empdir.=$entry->get_value('title').";
    $empdir.="</span><br>";
    ...
    $empdir.="</div>";
}
print STDOUT $empdir;
$msg=$ldap->unbind();
```

SpamAssassin and E-mail Integration

I designed an e-mail gateway for our company back in 2001, and it's still the system we use today. I wrote about it in a previous *Linux Journal* article in the December 2001 issue. The system has been modified tremendously since then, but it still operates in the same basic way. It's simply a store, scan and forward agent. Because this all takes place on our Linux server, our Windows users were unable to see or retrieve false positives or have any control over their SpamAssassin whitelists. We solved this by building a set of CGI scripts to let our users modify their SpamAssassin preferences file and release their false positives from the spam trap on their own, using the intranet as the interface.

Users launch the mail management scripts from their My Intranet section on our home page (Figure 2). They choose which day's mail they want to view from a drop-down box and click a button to activate the selfserv.cgi script. There is no user identity information passed to the script, because it will obtain that information from an SSC lookup. After we do the initial SSC lookup, we call the gettempinfo.pl script again to get the current user's e-mail address, like this:

```
##: Get this user's email address
open(GETEMPINFO,"-|","gettempinfo.pl",$cn,"mail");
my $searchstring=<GETEMPINFO>;
close(GETEMPINFO);
```

The \$searchstring variable then becomes the base of the regular expression we use to search the /spam directory for spam belonging to this user. As the mail attribute coming from Active Directory is something typed in by human hands, we must do another check to make sure we aren't falling victim to typos:

```
##: Make sure this email address is valid
unless($searchstring=~ /^[a-z]*\@domain\.com$/) {
    print STDOUT "Content-Type: text/plain\n";
    print STDOUT "Access Denied: Your identity on \
        the network can't be verified.\n";
    return(0);
}
```

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

If these checks are successful, the script responds by showing users the requested day's spam in a table format with a list of option links on the side of each item (Figure 3). Users then can use the option links to have the script release the spam, whitelist its sender, blacklist its sender, produce a SpamAssassin report or simply display it as plain text. The script looks up the user's SSC information each time it's called and before any action is performed so that it knows whether or not to allow that action. I won't get into more detail here, because the functions of this script consist mostly of just moving files around in response to users' requests. I do want to mention the whitelist and blacklist options though.

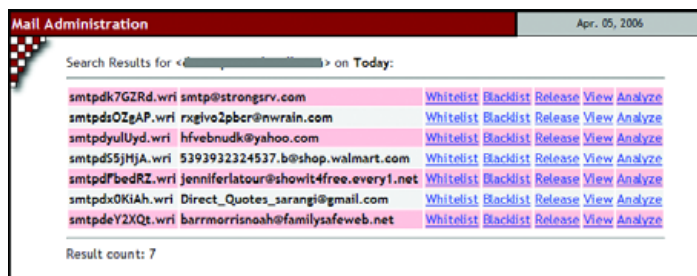


Figure 3. Options for Handling Trapped Spam

SpamAssassin holds its per-user configuration data in a file named `.spamassassin/user_prefs.cf` in each user's home directory. In a normal setup, where Linux is your main mail server, this is fine, but in our case, it won't work. Our Linux server is merely a scanning gateway that relays mail in and out, so it has no awareness of our users or their e-mail accounts. To solve this, we have to cheat a little. SpamAssassin's main configuration file is named `/etc/mail/spamassassin/local.cf`, and it reads this file every time it's started. It doesn't read only that file though. It actually reads all files in the `/etc/mail/spamassassin` directory that have a `.cf` extension. We can use this to our advantage and have our CGI script create files in this directory for each user's whitelist in a `$cn_prefs.cf` format. We have a cron job that restarts `spamd` every hour anyway to free memory, so this works out fine. The most important thing to remember if you use this method is that you have to do strict syntax checking to make sure users aren't whitelisting things like `*@hotmail.com` or using any other SpamAssassin directives. Even though these files have the appearance of private preference files to users, they actually are global to SpamAssassin, because they reside in the main config directory.

Microsoft SQL Server Integration

Our firm uses a time and billing system called CPAS. This software package holds all of our client and billing information as well as information used by our marketing manager to assemble mass mail-outs to our clients. We wanted to give our users access to this information to do some rudimentary data mining without having to contact administration every time. Because CPAS stores its information in a Microsoft SQL Server database, we had to use a piece of software called FreeTDS and the DBD::Sybase package from CPAN to interface to it from our Perl CGI scripts.

Four steps are involved in setting this up. The first thing to do is grab the latest FreeTDS package from the Internet and unpack the tarball. Next, `cd` into the unpacked directory, and execute the following commands:

```
> ./configure --prefix=/usr/local/freetds
```

```
> make
> su -c 'make install'
```

This sets up FreeTDS in its own directory, so it's easier for the Sybase module to find later. Next, we go into CPAN and get the DBD::Sybase package. Become root and execute the following commands:

```
> perl -MCPAN -e shell
> install DBD::Sybase
```

Feel free to force the install if some of the tests fail—that is pretty common according to the package's author. At this point, the software is installed, but we have to set up the FreeTDS configuration file. This file holds information about the databases to which you will be connecting. The configuration file is well documented, and you should be able to figure out the syntax easily. Here is a sample server entry:

```
[JACKSON5]
host = jackson5.domain.com
port = 1433
tds version = 4.2
```

Once FreeTDS is configured, you can access your database from your CGI scripts through the familiar DBI interface in Perl. Here is an example connection to a database called `concerts` running on a Windows server named JACKSON5:

```
#!/usr/bin/perl -w
use DBI; $ENV{'SYBASE'} = '/usr/local/freetds';
$dbh = DBI->connect('dbi:Sybase:server=JACKSON5', 'username', 'password')
or die 'connect';
$dbh->do("use concerts");
```

Notice that you have to put the location of your FreeTDS installation in an environment variable before you attempt a connection. The environment variable tells DBD::Sybase where to find the FreeTDS libraries. After that, you simply perform your queries as usual using DBI. If you are used to working with MySQL, I suggest you study up on the syntax used by Microsoft SQL Server. Some of it is very different from what you are used to.

Conclusion

I hope this article gives you some ideas and practical knowledge on how to better integrate your intranet with some of the more common systems found in a small business. An intranet shouldn't be only a news portal or electronic bulletin board. It should be an interactive tool for users and a time-saver for administrators. Users feel a level of comfort in a browser environment that they don't feel when searching through a filesystem or staring at a command line. Take advantage of that and your intranet will become a valuable asset to your business. ■

Resources for this article: www.linuxjournal.com/article/9270.

Dave Jones is the IT Manager at Pearce, Bevil, Leesburg & Moore in Birmingham, Alabama. He has been a network administrator for eight years. He spends his time blogging and writing software at www.sector62.com.

Add Web Porn Filtering and Other Content Filtering to Linux Desktops

How to set up the DansGuardian content filter with the lightweight Tinyproxy.

DONALD EMMACK

Microsoft users continue to adopt the Linux operating system and naturally expect to find content filters like the ones they used with Windows XP. Often, new Linux converts experiment on their standalone personal computers at home. Because many people object to some information and images readily found on the Internet, a content filtering system is top priority—especially because parents often share computers with kids, and constant adult supervision is not always possible.

Using DansGuardian with Tinyproxy is one way parents can supervise Internet content when they are away from the family computer. A versatile content filter, DansGuardian is open-source software for use in a noncommercial setting. If you want to use DansGuardian in a commercial setting, you can buy a license or buy SmoothGuardian. Working with DansGuardian is Tinyproxy, a small open-source program that understands and evaluates the information passing through the computer. Together they provide administrative controls to block objectionable content from the Internet.

Content Filtering at 5,000 Feet

DansGuardian is a collection of pass-through filters used to stop Internet Web pages with words, phrases and pictures you don't like or want others to see. The filters within DansGuardian act as an intermediary program between a client browser, like Firefox, and the Internet. Firefox makes the information request to DansGuardian. Then, DansGuardian passes the information to Tinyproxy, which communicates with the Internet.

Information coming back from the Internet passes through Tinyproxy and DansGuardian before it gets to the client browser. Only approved information gets through the filter and appears in the browser window. DansGuardian blocks restricted Web pages and replaces the unwanted content with an "access denied" security screen displayed in the browser window.

This has not been a high-level description of the filtering procedure. In fact, the way Tinyproxy and DansGuardian work together is complex and interesting. If you want to explore how this works, check out the DansGuardian "Flow of Events" page (see the on-line Resources). Here, you can find a more thorough discussion of filtering and how data passes between each program and the Internet.

What's important to know is you can define many words, phrases and specific locations you want DansGuardian to block. In addition to Web pages with text, DansGuardian also can filter pictures and prevent the downloading of certain files. This combination of filtering is superior to other methods that block access only to a list of banned sites.

With more than 20 different configuration files, setup of DansGuardian can appear complicated to new Linux users. However, the configuration files contain clear instructions on how to edit them for your needs. In my tests, I didn't need to make a lot of changes, because the default filtering arrangement is almost ideal for family use.

Installation

First, you need to install and configure DansGuardian and Tinyproxy. Second, it's important to adjust your desktop settings to prevent users from easily turning off content filtering.

Before installing, look through the package repository of your distribution to make sure it includes DansGuardian and Tinyproxy. The most simple way to install the programs is with a GUI package manager like Novel SUSE's YaST or Synaptic. For Debian, root users enter `apt-get install dansguardian tinyproxy`.

If you don't have these applications in your package repository, you can download DansGuardian and Tinyproxy from their respective Web pages (see Resources). After downloading, you will find generic installation instructions in the file named INSTALL.

Configuring DansGuardian and Tinyproxy

The next task is to customize configuration files for both Tinyproxy and DansGuardian. I use Ubuntu Dapper Drake for testing purposes, and so the directory and file illustrations are likely specific to this distribution. Other distributions organize files in a similar way; you just may need to look a little more to find the installation directory. For customizing features, the only tool necessary is a simple text editor, such as GNOME's gedit.

Using your text editor, as root user, open `/etc/dansguardian/dansguardian.conf`. Review the file and change `filterport`, `proxyip` and `proxyport` to match that shown below. Depending on your distribution, it also may be necessary to comment out the line starting with `UNCONFIGURED`:

```
# the port that DansGuardian listens to.
filterport = 8080

# the ip of the proxy--default is the loopback (this server)
proxyip = 127.0.0.1

# the port DansGuardian connects to proxy on
proxyport = 3128
```

DansGuardian generally connects to port 3128 by default, because that is the port used by the popular proxy called Squid. We can change this to the default port used by Tinyproxy (8888), or we can change the Tinyproxy port. In this case, we do the latter and change the port Tinyproxy uses to match the default Squid port.

For Tinyproxy, edit the file `/etc/tinyproxy/tinyproxy.conf` as root user. Look through this file, and make sure to change `User`, `Group`, `Port` and `ViaProxyName`, if necessary. The important thing to change is the port that Tinyproxy will use to match the DansGuardian connect port, which is 3128:

```
# Port to listen on.
#
Port 3128
```

Once you've finished with these changes, issue the command `tinyproxy` in your terminal, or if Ubuntu-based, type `sudo /etc/init.d/tinyproxy start`. This starts the proxy, and you're now ready to finish off the installation by adjusting your browser preferences. If you want to learn more, look at the DansGuardian documentation links (see Resources) for a description of this process.

Adjust Your Browser Settings

Ubuntu comes with Firefox as the preferred client browser, so the instructions here are specific to Firefox. Other client browsers will likely have similar capabilities and documentation to show how to mimic

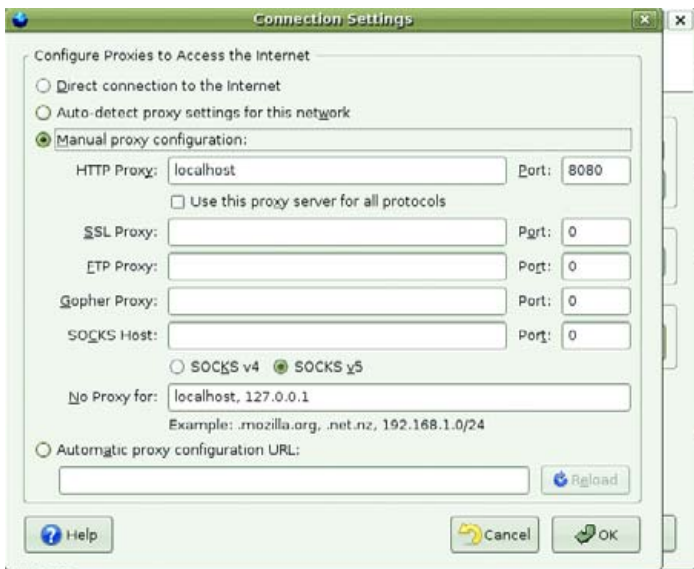


Figure 1. Set up your browser to use the proxy.



Figure 2. A Typical DansGuardian Access Denied Page

these instructions.

This last installation step points the browser at port 8080, so it sends data only through DansGuardian and Tinyproxy. With Firefox, go to Edit→Preferences→General tab→Connection Settings to see the screen shown in Figure 1. As shown, select manual proxy configuration, enter localhost and port 8080. This assumes you are going to install and use DansGuardian and Tinyproxy on every workstation. If you set up DansGuardian and Tinyproxy on a separate server, then you need to enter the name or IP address of the server machine that runs DansGuardian and Tinyproxy instead of the word localhost in the HTTP Proxy line.

Restart your browser and test how well the filter works.

When testing the new filter, you should see an access denied screen similar to the one shown in Figure 2. Before going any further, it's a good idea to look for problems you may find with the default filter settings. For example, I often download .tar and other executable files. The default configuration file stops these files from download. To fix this problem, you need to edit the `bannedextensionlist.txt` file, and place a # to comment out the file extensions you want to let through the filter.

To be thorough, you should look through all default configuration .txt files with DansGuardian to tailor how you want the filters to react. You won't know all the situations you'll run into at first, but this is a good opportunity to gain an understanding of this application's powerful features.

Some Vulnerabilities

No system is perfect, and there are several obvious ways to defeat DansGuardian and Tinyproxy. The most noteworthy is how easily users can bypass the proxy and filters. Without further protection, a user can restore Firefox's preferences back to Direct Connection, which bypasses DansGuardian and Tinyproxy. Once reversed, users have unrestricted access to the Internet.

However, there are more ways to secure the DansGuardian filters

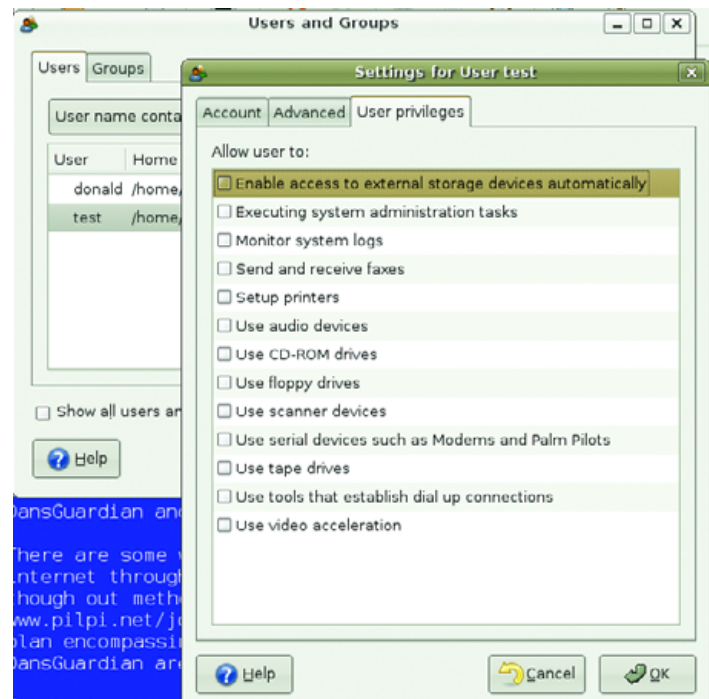


Figure 3. Ubuntu Dapper Drake User Privilege Settings

further by forcing all communication with the Internet through port 8080. A link on the DansGuardian documentation Web page explains a well-thought-out method of using FireHol to force this condition on all Internet thoroughfares (see Resources).

For the novice user, an easier approach is to set up a filtering plan that includes restricted user privileges, locked browser preferences and making sure the proxy filters start each time the computer reboots.

For test purposes, I created a new user account on Ubuntu Dapper Drake (Figure 3). Using the checklist features, I severely limited the capability of the user test. Although these privileges could be just right for anyone who has no computer experience or who is plainly not trustworthy. Utilities like update-rc.d and fconfd define certain programs to start at the system boot. I used a bootup manager called BUM to make DansGuardian and Tinyproxy start at each boot.

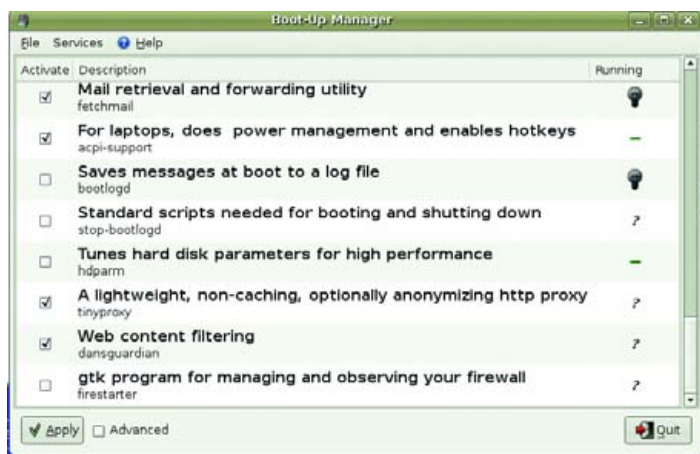


Figure 4. Set up DansGuardian and Tinyproxy to run every time you boot Linux.

Finally, I decided to lock down the preferences of Firefox. Restricting Firefox's preferences is not as difficult as it may sound. An older copyrighted article titled "HOWTO Lock Down Mozilla Preferences for LTSP" by Warren Togami (see Resources) describes how to carry this out in great detail. Although, I didn't want to mess with byte shift coding to achieve similar results.

After rummaging through Mozilla.org's Web site, I chose to add lockPref statements to my Firefox configuration file to keep users from changing connection settings. I edited the file /usr/lib/firefox/firefox.cfg to appear as the one shown in Figure 5. The last three lines force a manual proxy selection on localhost, port 8080. After saving this file and restarting Firefox, you can't reset the connection settings. Further, other users without administrative privileges could not quickly change the settings and bypass the filters.




Figure 5. Lock down Firefox settings so they can't be changed without administrative privileges.

Maintenance

After customizing the filters to your liking, it's important to realize that some settings become stale. Blacklisted sites and new phrases

are likely to go out of date sooner than others. New Web sites you will want to block come on-line often, and new word combinations can make past phrases obsolete. Looking through the Extras link on the DansGuardian site, you will find more information on blacklists. In addition, several users have contributed scripts to automate blacklist generation and update.

As an alternative, URLblacklist.com allows new users to download their first file free. Afterwards, you can sign up for a periodic subscription for access to the latest-and-greatest information. Instructions for applying the new data for DansGuardian are on the Web site.

Another consideration is whether the proxy and filter will slow down Internet surfing and page loading. Some users will suffer a small impact on Web surfing performance when using Tinyproxy. In my own testing, I noticed a slight delay, plus a couple of issues with my browser cache. Clearing the Firefox cache with Ctrl-Shift-Del fixed the cache problems right away. Occasionally, it has been necessary to restart Tinyproxy. After doing so, my Internet performance improved. Although annoying at times, these small issues are acceptable trade-offs.

Log File Review

Both DansGuardian and Tinyproxy make log files for administrators to review. Within /var/log, you should find directories for DansGuardian and Tinyproxy. Using an editor, open the files and search through the data to find out what's been happening on the computer. Sequentially stored data and clear comment fields make the file easier to understand. For DansGuardian, there is a user-contributed add-on script for searching and displaying the results in a more user-friendly format.

One feature not found in DansGuardian is the capacity to e-mail the log files to a third party for review. This can be a real deterrent for some people if they know they have an accountability partner watching their actions on the Internet.

Some Final Thoughts

Before settling on this solution for content filtering, consider what your overall requirements are in the upcoming months. If you have only one computer to deal with and you don't mind tinkering with configuration files, DansGuardian is probably a good choice. Alternatively, SmoothGuardian looks like a great buy for \$90 US. Plus, the software includes a user-friendly Web-based interface and nontechnical installation.

Nevertheless, setup of DansGuardian and Tinyproxy is well within the scope of new Linux users, and the free price fits most budgets nicely. Using this article and its references as a guide, you shouldn't have too much difficulty getting up and running. Even if you do battle a few problems, using Google to search for answers is easy. Plus, there is also a Web content filtering portal linked to the DansGuardian home page (see Resources) and an IRC chat location.


Overall, DansGuardian and Tinyproxy are frontrunners in the Open Source world and help ease the transition from the Microsoft Windows environment. I think you'll find flexible filtering and lightweight proxy overhead make this a good combination for small networking environments. ■

Resources for this article: www.linuxjournal.com/article/9291.

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

Come Together

Unique innovations are wonderful, but do Linux distributions have to differentiate at such low levels?



Nick Petreley, Editor in Chief

The whole PC world is plagued by a lack of good standards. Some of the most frustrating standards problems are hardware-related. For example, what brainiac thought it was a good idea to make the FireWire connectors and USB connectors on motherboards identical? The motherboard manuals are usually careful to point out that if you mix these up, you can damage the motherboard. That's nice, but who made it possible to mix them up in the first place? Dumb.

It's just as troubling to see a continuing lack of good, comprehensive standards among Linux distributions. As with hardware, you can almost always find a way to make something work if you are careful and know what you're doing. But that's no excuse for the lack of standards across distributions, and the few inadequate standards that exist.

Here's what inspired this complaint. If you've been following my columns, you'll know that I've been trying to put together a MythTV box. I followed several how-to pages for installing special drivers for the tuner cards I have tried. Most of the published instructions, including those linked to

by some hardware vendors, tell you to place firmware everywhere but the place Ubuntu stores firmware. Ubuntu looks for firmware in the `/lib/firmware/<kernel version>` directory. Most instructions tell you to put the firmware in `/usr/lib/hotplug/firmware`. One card, the Hauppauge PVR-150/500, wants firmware files in multiple locations, including the `/lib/modules/` directory. It uses different filenames depending on the version of the kernel and driver. I've tested three cards so far, and I finally ran out of patience and used a shotgun approach. I put copies of the firmware just about everywhere but my son's sock drawer. All the drivers work now. I have no idea which copies of the firmware files they are finding, but I don't care anymore.

Personally, I like the Ubuntu approach to locating firmware. Ubuntu uses `udev`, which many agree is superior to `hotplug`. It lets you install separate versions of firmware based on the version of the kernel.

Some may argue that this differentiation is what open source is all about. If Ubuntu's choice is good enough, other distributions will cream-skim it, and it will become the standard. Fair enough, but wouldn't it be more efficient for customers if the distributors simply agreed on such fundamentals as `udev` and where to put firmware? At least that way we'd be less likely to run across how-to pages that don't apply to our chosen distribution.

As much as I like this one thing about Ubuntu, Ubuntu is far from perfect when it comes to establishing or observing standards. Try to install a vanilla kernel on Ubuntu and see for yourself. You'll notice that you can no longer mount some disk partitions. Ubuntu, by default, installs and uses a logical volume manager (LVM) and enterprise volume management services (EVMS), one or both of which break how Ubuntu works if you use a vanilla kernel. I managed to fix the mount problem by editing the configuration files for LVM and EVMS to ignore all the drives on my system. The next version of Ubuntu will add

ivman, yet another volume manager. I can't wait to find out what I'll have to reconfigure when the new Ubuntu is ready.

Unfortunately, my suggestion that distributors collaborate is utopian and unrealistic. They don't even work as a team in ways that would benefit them most, such as pressuring hardware vendors to preload Linux. When it comes to standards, most distributors aren't even willing to agree on a package format let alone build a package system where you could install a Mandriva RPM in Fedora without running into dependency problems. They can't agree on where to put firmware files or whether EVMS should be part of the basic system.

The best possible solution would be for all major distributors to build on a single base distribution. This was one of the original ideas posed when Linux Standard Base was first formed, but distributors rejected the idea in spite of the fact that it would save them all a lot of duplicated effort. Why are distributors disinclined to agree on a comprehensive standard distribution? Competition. A standard base distribution would lower the barrier of entry for new competing distributions. Put more bluntly, despite all the lip service Linux distributors give to how their commitment to open source and freedom empowers end users, they really do like having a degree of customer lock-in. Their lock-in just isn't as severe, obvious, destructive or effective as Microsoft's lock-in.

Don't get me wrong. I don't want to see the Linux market homogenized so much that distributions start to disappear. I'm glad there are many distributions from which to choose. I would simply like to see them differentiate their distributions at a much higher level, a level that eliminates needless compatibility problems. But I confess that there are times when frustration leads me to the temptation to start a crusade to get everyone to run Debian. What do you think? ■

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



Rackspace – Managed Hosting Backed by Fanatical Support™

Fast servers, secure data centers and maximum bandwidth are all well and good. In fact, we invest a lot of money in them every year. But we believe hosting enterprise class web sites and web applications takes more than technology. It takes Fanatical Support.

Fanatical Support isn't a clever slogan, but the day to day reality our customers experience working with us. It's how we have reimagined customer service to bring unprecedented responsiveness and value to everything we do for our customers. It starts the first time you talk with us. And it never ends.

Contact us to see how Fanatical Support works for you.

1.888.571.8976 or visit www.rackspace.com



Thanks for honoring us with the
2005 Linux Journal Readers' Choice Award for
"Favorite Web-Hosting Service"



The World's Fastest InfiniBand Switch

From a Company You've Trusted for 24 Years

Microway's **FasTree™** DDR InfiniBand switches run at 5GHz, twice as fast as the competition's SDR models. FasTree's non-blocking, flow-through architecture makes it possible to create 24 to 72 port modular fabrics which have lower latency than monolithic switches. They aggregate data modulo 24 instead of 12, improving nearest neighbor latency in fine grain problems and doubling the size of the largest three hop fat tree that can be built, from 288 to 576 ports. Larger fabrics can be created linking 576 port domains together.



▲ 72 Port FasTree™ Configuration

Working with QLogic's InfiniPath InfiniBand Adapters, the number of hops required to move MPI messages between nodes is reduced, improving latency. The modular design makes them useful for SDR, DDR and future QDR InfiniBand fabrics, greatly extending their useful life. Please send email to fastree@microway.com to request our white paper entitled *Low Latency Modular Switches for InfiniBand*.

Harness the power of 16 Opteron™ cores and 128 GB in 4U

Microway's **QuadPuter®** includes four or eight AMD dual core Opteron™ processors, 1350 Watt redundant power supply, and up to 8 redundant, hot swap hard drives—all in 4U. Dual core enables users to increase computing capacity without increasing power requirements, thereby providing the best performance per watt. Constructed with stainless steel, QuadPuter's RuggedRack™ architecture is designed to keep the processors and memory running cool and efficiently. Hard drives are cooled with external air and are front-mounted along with the power supply for easy access and removal. The RuggedRack™ with an 8-way motherboard, 8 drives, and up to 128 GB of memory is an excellent platform for power- and memory-hungry SMP applications.



Call us first at 508-746-7341 for quotes on clusters and storage solutions. Find testimonials and a list of satisfied customers at microway.com.

◀ QuadPuter® Navion™ with hot swap, redundant power and hard drives and four or eight dual core Opterons, offering the perfect balance between performance and density



508.746.7341 microway.com