

Rocks | Lightning Hacks | Dojo | Nagios | VirtualBox | Privacy

LINUX JOURNAL


Since 1994: The Original Magazine of the Linux Community

JUNE 2009 | ISSUE 182 | www.linuxjournal.com

Build a
High-Performance
Compute Cluster
with **Rocks**

Using **VirtualBox**,
JumpBox and
Virtual Appliances

AND THE WINNERS ARE...



Your Favorite Gadgets, Programs, Tools, Hardware and More!


**How-To:
Dojo's Grid
Widget**

**Tools for
Validating
HTML**

**Monitoring
Dell Servers
with Nagios**



Embedded Wireless Dream Machines.




RouterStation

Featuring a fast 680MHz MIPS 24K CPU, 64MB RAM, and 16MB Flash; RouterStation provides an excellent horsepower for a variety of processor intensive multi-radio system applications.

Up to 3 mini-PCI radios, 3 10/100 ethernet interfaces, a 5A power supply for multiple hi-power card support, USB 2.0, and enhanced temperature operating performance and ethernet ESD protection for carrier applications.

MSRP \$59



RouterStation Pro

In response to the outstanding demand for our initial RouterStation OEM platform, Ubiquiti Networks announces the RouterStation Pro. Breakthrough Price/Performance with a \$79 USD MSRP.


Pro Version Enhancements:

- 48V 802.3af Power Over Ethernet
- 4-Port Gigabit Ethernet Switch
- 256MB RAM
- On Board SDIO Support
- On Board, USB 2.0, RS232/dB9, and DC power jacks

MSRP \$79

YOUR HIGH PERFORMANCE COMPUTING HAS ARRIVED.

The ServersDirect® Systems with the Intel® Xeon® Processor helps you simplify computing operations, accelerate performance and accomplish more in less time



STARTING AT **\$899**

ENTRY LEVEL INTELLIGENT SERVER

SDR-S1341-T00 is among our most cost-effective 1U Xeon Servers, and it is ideal for large high-performance computing deployments



STARTING AT **\$959**

APPLICATION SERVER

Refresh your servers with new **SDR-S1337-T02** powered by Intel® Xeon® processor 5500 series, based on intelligent performance, automated energy efficiency and flexible virtualization.



SDR-S1343-T04
STARTING AT **\$1,099**

1U INTEL® XEON® PROCESSORS 5500 SERIES SERVER W/ 4X 3.5" HOT-SWAP SATA DRIVE BAYS

- Supermicro 1U Rackmount Server with 560W Power Supply
- Supermicro Server Board w/Intel® 5520 Chipset
- Support up to Dual Intel® 5500 series Xeon® Quad/Dual-Core, with QPI up to 6.4 GT/s
- Support up to 96GB DDR3 1333/ 1066/ 800MHz ECC Reg.DIMM
- 4x 3.5" Hot-swap SATA Drive Bays
- Intel® 82576 Dual-Port Gigabit Ethernet Controller



SDR-S2311-T08
STARTING AT **\$1,159**

2U INTEL® XEON® PROCESSORS 5500 SERIES SERVER W/ 8X 3.5" HOT-SWAP SAS/SATA BAYS

- Supermicro 2U Rackmount Server with 560W Power Supply
- Supermicro Server Board w/Intel® 5500 Chipset
- Support up to Dual Intel® 5500 series Xeon® Quad/Dual-Core, with QPI up to 6.4 GT/s
- Support up to 24GB DDR3 1333/ 1066/ 800MHz ECC Reg.DIMM
- 8x 3.5" Hot-swap SATA Drive Bays
- Dual Intel® 82574L Gigabit Ethernet Controller



SDP-IP308-T10
STARTING AT **\$1,599**

PEDESTAL INTEL® XEON® PROCESSORS 5500 SERIES SERVER W/ 10X HOT-SWAP (OPT.) SATA BAYS

- Intel Pedestal Chassis w/ 750W (1+1) Power Supply
- Supermicro Server Board w/Intel® 5520 Chipsets
- Support up to Dual Intel® 5500 series Xeon® Quad/Dual-Core, with QPI up to 6.4 GT/s
- Support up to 96GB DDR3 1333/ 1066/ 800MHz ECC Reg./unbuffered DIMM
- Option 10x 3.5" Hot-swap SATA Bays
- Intel® 8257EB Dual-port Gigabit Ethernet Controller



SDR-S4313-T24
STARTING AT **\$1,899**

4U INTEL® XEON® PROCESSORS 5500 SERIES SERVER W/ 24X 3.5" HOT-SWAP SAS/SATA BAYS

- Supermicro 4U Rackmount 900W (1+1) Red. Power Supply
- Supermicro Server Board w/ Dual Intel® 5520 Chipsets
- Support up to Dual Intel® 5500 series Xeon® Quad/Dual-Core, with QPI up to 6.4 GT/s
- Support up to 144GB DDR3 1333/ 1066/ 800MHz ECC Reg. DIMM
- 24x 3.5" Hot-swap SATA Drive Bay
- Intel® 82576 Dual-port Gigabit Ethernet Controller



SDR-S3305-T16
STARTING AT **\$1,979**

3U INTEL® XEON® PROCESSORS 5500 SERIES SERVER W/ 16X 3.5" HOT-SWAP SAS/SATA BAYS

- 3U Rackmount Server with 1+1 900W Red. Power Supply
- Supermicro Server Board w/ Dual Intel® 5520 Chipsets
- Support up to Dual Intel® 5500 series Xeon® Quad/Dual-Core, with QPI up to 6.4 GT/s
- Support up to 96GB DDR3 1333/ 1066/ 800MHz ECC Reg.DIMM
- 16x Hot-swap SAS/SATA Drive Bays
- Intel® Dual 82576 Dual-Port Gigabit Ethernet (4 ports)



SDR-C9303-T50
STARTING AT **\$4,339**

9U INTEL® XEON® PROCESSORS 5500 NEHALEM SERIES SERVER W/ 50X HOT-SWAP SATA II / SAS BAYS

- 9U Chassis with 1620W Redundant Power Supply
- Supermicro Server Board w/ Dual Intel® 5520 Chipsets
- Support up to Dual Intel® 5500 series Xeon® Quad/Dual-Core, with QPI up to 6.4 GT/s
- Support up to 144GB DDR3 1333/ 1066/ 800MHz ECC Reg. DIMM
- 50 x 3.5" Internal SATA Drives Trays
- Intel® 82576 Dual-port Gigabit Ethernet Controller



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

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

1.877.727.7886 / www.ServersDirect.com

Intel, Intel logo, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, Pentium, and Pentium III Xeon are trademarks of Intel Corporation or it's subsidiaries in the United States and other countries.



CONTENTS

JUNE 2009

Issue 182



FEATURE

44 READERS' CHOICE AWARDS 2009

James Gray

ON THE COVER

- Build a High-Performance Compute Cluster with Rocks, p. 56
- Using VirtualBox, JumpBox and Virtual Appliances, p. 26
- And the Winners Are..., p. 44
- How-To: Dojo's Grid Widget, p. 66
- Tools for Validating HTML, p. 20
- Monitoring Dell Servers with Nagios, p. 75

The Best Technical Training for SharePoint Comes to Boston!

Attend



June 22-24, 2009 → Boston

Hyatt Regency Cambridge
Cambridge, MA



Six Great Reasons to Attend SPTechCon Boston

1. There's a shortage of SharePoint experts! Develop your skills, and improve your own professional standing!
2. Learn from the brightest minds in the SharePoint universe! Most of our speakers either are Microsoft engineers or have achieved MVP status based on their in-depth knowledge of SharePoint.
3. Find the best third-party tools and meet informally with the experts in our exhibit hall!
4. Take only the classes that work best for you. With more than 60 classes and workshops to choose from, you can make SPTechCon your own!
5. Bring a group of developers, IT pros and business teams to improve your whole organization's skills – and get a discount to boot! Contact us for group registration discounts!
6. The technical classes and workshops at SPTechCon are focused on practical techniques and practices you can put to work today!

 **REGISTER** early to take advantage of our **Early Bird Rates!**

For more information, go to www.sptechcon.com



PRODUCED BY
BZ Media
SYSTEMS MANAGEMENT NEWS
SDTimes
The Industry Perspective for Software Development Managers

CONTENTS

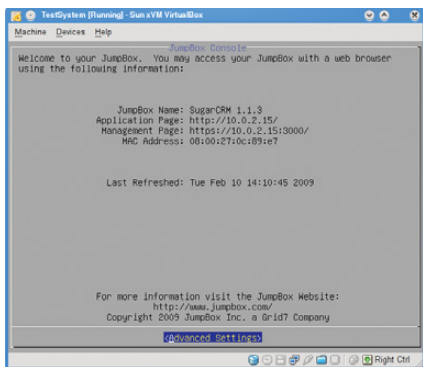
JUNE 2009

Issue 182

COLUMNS

20 REUVEN M. LERNER'S AT THE FORGE
Checking Your HTML

26 MARCEL GAGNÉ'S COOKING WITH LINUX
Serious Cool, Sysadmin Style!



32 DAVE TAYLOR'S WORK THE SHELL
Deal or No Deal!

34 KYLE RANKIN'S HACK AND /
Lightning Hacks Strike Twice

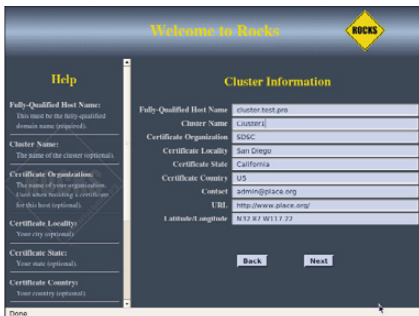
80 DOC SEARLS' EOF
The Post-Monopoly Game

IN EVERY ISSUE

- 10** CURRENT_ISSUE.TAR.GZ
- 12** LETTERS
- 16** UPFRONT
- 38** NEW PRODUCTS
- 40** NEW PROJECTS
- 65** ADVERTISERS INDEX
- 79** MARKETPLACE

HIGH-PERFORMANCE COMPUTING

56 BUILDING A LINUX-BASED HIGH-PERFORMANCE COMPUTE CLUSTER
Don't toss 'em, cluster 'em.
Tom Lehmann

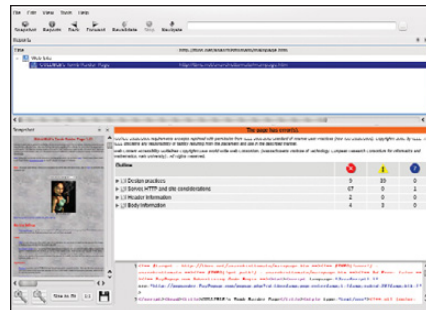


INDEPTH

66 DOJO'S INDUSTRIAL-STRENGTH GRID WIDGET
Display arbitrarily large volumes of tabular data.
Matthew Russell

ID	Label
554036	qux
554037	baz
554038	baz
554039	foo
554040	qux
554041	bar

75 SNMP MONITORING WITH NAGIOS
Monitor Dell servers with SNMP and Nagios.
Jason Ellison



40 RIGHTWEBPAGE

Next Month

MOBILE LINUX

Linux is on the move, and we mean physically on the move. So unless you're using two tin cans and string to communicate, chances are, something you're carrying has Linux in it. Phones, Netbooks, you name it—if it doesn't run Linux now, it probably will soon.

Tired of just playing music on the go? How about composing it on the go with the portable Indamixx music studio? Need a server when you get there? How about turning your Western Digital MyBook into a portable server? Need something for some light reading? How about the new Kindle 2 from Amazon? Having trouble with that small screen? Try the Conkeror Web browser (a must-have if you're an Emacs fan).

And we can't do Mobile if we don't do Android. Find out how to get started programming on Android with its high-level Java API. Now you too can write telephony apps.

USPS LINUX JOURNAL (ISSN 1075-3583) (USPS 12854) is published monthly by Belltown Media, Inc., 2211 Norfolk, Ste 514, Houston, TX 77098 USA. Periodicals postage paid at Houston, Texas and at additional mailing offices. Cover price is \$5.99 US. Subscription rate is \$29.50/year in the United States, \$39.50 in Canada and Mexico, \$69.50 elsewhere. POSTMASTER: Please send address changes to Linux Journal, PO Box 16476, North Hollywood, CA 91615. Subscriptions start with the next issue. Canada Post: Publications Mail Agreement #41549519. Canada Returns to be sent to Bleuchip International, P.O. Box 25542, London, ON N6C 6B2

Polywell Linux Solutions

More Choices, Excellent Service, Great Value!
Serving the Industry for More Than 20 Years

Netdisk 8000V Quiet Performance NAS Storage



4TB **\$1,399**
8TB **\$2,399**
12TB **\$2,999**

- Dual Gigabit LAN
- RAID-5, 6, 0, 1, 10
- Hot Swap, Hot Spare
- Linux, Windows, Mac
- E-mail Notification
- Tower or Rackmount



Silent Eco Green PC

NVIDIA® nForce Chipset and GeForce Graphics
Energy efficient, Quiet and Low Voltage Platform. starts at **\$199**



The Best Terminal PC

Fanless Silent ITX PC

1G DDR2, Solid State Drive starts at **\$299**
Low-Voltage processor, Low-profile Add-on Available
Excellent for Linux Appliance



4U 24Bay 36TB Storage Server

Hardware RAID-6, NAS/iSCSI/SAN Storage
Mix SAS and SATA, 4 x GigaLAN or 10Gbit LAN



Mini-1U Server for Data Center ISP

Dual-Core or Quad-Core Processor
4GB to 8GB RAM, 2 x 500GB RAID HD
Linux Server Starts at **\$499**



Polywell OEM Services, Your Virtual Manufacturer
Prototype Development with Linux/FreeBSD Support
Small Scale to Mass Production Manufacturing
Fulfillment, Shipping and RMA Repairs

- 20 Years of Customer Satisfaction
- 5-Year Warranty, Industry's Longest
- First Class Customer Service

888.765.9686

linuxsales@polywell.com

www.polywell.com/us/Lx



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

NVIDIA, nForce, GeForce and combinations thereof are trademarks of NVIDIA Corporation. Other names are for informational purposes only and may be trademarks of their respective owners.

LINUX JOURNAL™

Since 1994: The Original Magazine of the Linux Community

Digital Edition Now Available!

Read it first

Get the latest issue before it
hits the newsstand

Keyword searchable

Find a topic or name
in seconds

Paperless archives


Download to your computer for
convenient offline reading

Same great magazine

Read each issue in
high-quality PDF

Try a Sample Issue!

www.linuxjournal.com/digital



LINUX JOURNAL

Executive Editor Jill Franklin
jill@linuxjournal.com

Senior Editor Doc Searls
doc@linuxjournal.com

Associate Editor Shawn Powers
shawn@linuxjournal.com

Associate Editor Mitch Frazier
mitch@linuxjournal.com

Art Director Garrick Antikajian
garrick@linuxjournal.com

Products Editor James Gray
newproducts@linuxjournal.com

Editor Emeritus Don Marti
dmarti@linuxjournal.com

Technical Editor Michael Baxter
mab@cruzio.com

Senior Columnist Reuven Lerner
reuven@lerner.co.il

Chef Français Marcel Gagné
maggagne@salmar.com

Security Editor Mick Bauer
mick@visi.com

Hack Editor Kyle Rankin
lj@greenfly.net

Contributing Editors

David A. Bandel • Ibrahim Haddad • Robert Love • Zack Brown • Dave Phillips • Marco Fioretti
Ludovic Marcotte • Paul Barry • Paul McKenney • Dave Taylor • Dirk Elmendorf

Proofreader Geri Gale

Publisher Carlie Fairchild
publisher@linuxjournal.com

General Manager Rebecca Cassity
rebecca@linuxjournal.com

Sales Manager Joseph Krack
joseph@linuxjournal.com

Sales and Marketing Coordinator Tracy Manford
tracy@linuxjournal.com

Associate Publisher Mark Irgang
mark@linuxjournal.com

Webmistress Katherine Druckman
webmistress@linuxjournal.com

Accountant Candy Beauchamp
acct@linuxjournal.com

Linux Journal is published by, and is a registered trade name of, Belltown Media, Inc.
PO Box 980985, Houston, TX 77098 USA

Reader Advisory Panel

Brad Abram Baillo • Nick Baronian • Hari Boukis • Caleb S. Cullen • Steve Case
Kalyana Krishna Chadalavada • Keir Davis • Adam M. Dutko • Michael Eager • Nick Faltys • Ken Firestone
Dennis Franklin Frey • Victor Gregorio • Kristian Erik • Hermansen • Philip Jacob • Jay Kruiuzenga
David A. Lane • Steve Marquez • Dave McAllister • Craig Oda • Rob Orsini • Jeffrey D. Parent
Wayne D. Powell • Shawn Powers • Mike Roberts • Dracron Smith • Chris D. Stark • Patrick Swartz

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
Bruce Perens
Bdale Garbee, Linux CTO, HP
Danese Cooper, Open Source Diva, Intel Corporation

Advertising

E-MAIL: ads@linuxjournal.com
URL: www.linuxjournal.com/advertising
PHONE: +1 713-344-1956 ext. 2

Subscriptions

E-MAIL: subs@linuxjournal.com
URL: www.linuxjournal.com/subscribe
PHONE: +1 818-487-2089
FAX: +1 818-487-4550
TOLL-FREE: 1-888-66-LINUX
MAIL: PO Box 16476, North Hollywood, CA 91615-9911 USA
Please allow 4-6 weeks for processing address changes and orders
PRINTED IN USA


LINUX is a registered trademark of Linus Torvalds.



Your Applications Will Run Faster With Next Generation Microway Solutions!

TriCom™ X

- QDR/DDR InfiniBand HCA
- ConnectX™ Technology
- 1µsec Latency
- Switchless Serial Console
- NodeWatch™ Remote Management



Teraflop GPU Computing

For Workstations and HPC Clusters

- NVIDIA® Tesla™ GPU with 240 Cores on One Chip
 - CUDA™ SDK
- NVIDIA® Quadro® Professional Graphics
- AMD® FireStream™ GPU
 - Stream SDK with Brook+



NumberSmasher®

Large Memory Scalable SMP Server

- Scales to 1 TB of Virtual Shared Memory
- Up to 128 CPU Cores
- 8U System Includes 32 Quad Core CPUs
- QDR 1 µsec Backplane

FaSTree™ X

- Mellanox® InfiniScale™ IV Technology
- QDR/DDR InfiniBand Switches
- Modular Design
- 4 GB/sec Bandwidth per Port
- QSFP Interconnects
- InfiniScope™ Real Time Diagnostics



Call the HPC Experts at Microway to Design Your Next
High-Reliability Linux Cluster or InfiniBand Fabric.

508-746-7341

Sign up for Microway's
Newsletter at
www.microway.com

**Microway**
Technology you can count onsm

Everything you need for

Limited Time Offer:
Get 50% off for the first
6 months and so much
more when you sign up
for a 1 year plan!
Visit www.1and1.com
for details!

Now
50%
off
Offer extended



1-877-GO-1AND1

a successful website

DOMAINS

NEW Savings!

Reserve your unique web address today and save.

Protect your personal information. Private domain registration, a 2 GB e-mail account, domain forwarding and a starter website-building tool are included for FREE!

.info
\$0.99
FIRST YEAR*

.com .net
.org .us

~~\$8.99/year~~
\$6.99
FIRST YEAR*

Now on sale!*

WEB HOSTING

Everything you need for a professional website.

1&1 WebsiteBuilder, 1&1 E-mail Marketing Tool, Mailing List, Driving Directions, Photo Gallery, 1&1 Blog and more! Search advertising vouchers are included with all plans!

FREE Domain Names!



1&1 BUSINESS PACKAGE

- 3 domain names
- 2,500 e-mail accounts
- 250 GB web space/ 2,500 GB monthly transfer volume

~~\$9.99/month~~
\$4.99
PER MONTH*
(first 6 months)

Now 50% off
Offer extended

PREMIUM SERVERS

Designed specifically for high performance needs.

Our premium servers feature top-of-the-line AMD Opteron™ processors with energy efficient technology, reducing costs and environmental impact with increased performance-per-watt. 1&1 matches 100% of the energy consumed in our data center with Renewable Energy Certificates.



1&1 BUSINESS SERVER II

- Dual-Core Opteron™ Processor 1218
- 2 x 2.6 GHz, 4 GB DDR RAM,
- 2 x 500 GB hard drive
- RAID 1 included for FREE!

~~\$199.99/month~~
\$99.99
PER MONTH*
(first 6 months)

Now 50% off
Offer extended

Now accepting
PayPal™



*Offers valid for a limited time only. .info Special offer valid as of May 4, 2009. Private domain registration not available with .us domains. Server prices based on Linux servers. Setup fee, minimum contract term, and other terms and conditions may apply. Visit www.1and1.com for full promotional offer details. Program and pricing specifications, availability and prices subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet AG, all other trademarks are the property of their respective owners. © 2009 1&1 Internet, Inc. All rights reserved.

www.1and1.com

1&1



SHAWN POWERS

The Issue You Write

That's right, it's our annual Readers' Choice issue of *Linux Journal*. Every year, we ask you, our readers, to share what you think really stood out over the past year. And, then we tell you about it. Certainly, it seems like an easy way for us to avoid writing articles for you, but rest assured, the Readers' Choice issue isn't our excuse to take a cruise for a month and leave you to your own wiles. It's more like a community issue. You give, we give, and everyone is happy. Well, maybe not everyone. Emacs users for instance, or perhaps KDE users. But, I won't ruin the results for you.

In fact, with the Readers' Choice theme this month, it allows our columnists a little more flexibility in regard to what they write, and they didn't disappoint. Marcel Gagné shows us a handful of ways to install and test products that you may have never considered before. From virtual machines to jumpboxes, if you want to try out some server applications, you'll want to read this month's *Cooking with Linux* column.

Kyle Rankin has another go at *Lightning Hacks* this month and gives us four quick but useful tips that make life a bit easier behind the keyboard. Kyle claims he got the idea for *Lightning Hacks* from the common *Lightning Talks* featured at many conferences. I suspect he's just jealous of our daily video tech tips over at LinuxJournal.com. Sadly, I don't have any way to validate my claim. Speaking of validation (nice segue, no?), Reuven M. Lerner shows us how to validate HTML code. Linux users are big fans of open standards. Unfortunately, we sometimes fail to follow them ourselves. Reuven aims to change that this month, so be sure to read his column to learn more.

Last year, we did an issue dedicated to high-performance computing. We got a lot of positive feedback from that issue and thought the Readers' Choice issue would be a great

place to put in a High Performance section. Tom Lehmann demonstrates how easy it is to create your own computer cluster. We met Tom at the Supercomputing conference and asked him to prove it was easy to set up a cluster. I think I owe him a soda or something, because this month, he shows us step by step how to set up a Rocks cluster of our very own.

On the programmer's side, Matthew Russell is back this month showing off Dojo's Grid Widget. Sometimes displaying large amounts of data is difficult, but Dojo makes it a bit easier. If system administration is your thing, be sure to check out Jason Ellison's article on SNMP monitoring with Nagios. There's no such thing as too much monitoring data, and even if there were, Dojo can help us display it!

To finish up the Readers' Choice issue nicely, we have Dave Taylor's script-fu to help figure out the odds in the game show, *Deal or No Deal*. I don't think it's quite as frowned upon as counting cards in Vegas, but I'm not sure Howie Mandel would look too kindly on a contestant with a laptop calculating odds—perhaps if the script could be ported to an Android handset....

In the end, this issue and every issue is all about you, the reader. Enjoy the Readers' Choice issue. If your tastes line up with the majority of voters, you can bask in the comfort of commonality. If your application of choice didn't even make the list, you can smugly assure yourself the rest of the world just isn't as enlightened. After all, as Linux users, we're used to going against the grain. It has worked for us so far; I see no reason to change now. ■

Shawn Powers is the Associate Editor for *Linux Journal*. He's also the Gadget Guy for LinuxJournal.com, and he has an interesting collection of vintage Garfield coffee mugs. Don't let his silly hairdo fool you, he's a pretty ordinary guy and can be reached via e-mail at shawn@linuxjournal.com. Or, swing by the [#linuxjournal](https://freenode.net) IRC channel on [Freenode.net](https://freenode.net).



Systems

Dependable, Cost Effective Linux-Based Server Solutions

Designed for professional-grade performance, cost-effective ZT Reliant servers and desktops maximize your productivity and optimize your IT budget. An established USA-based manufacturer celebrating our 15th anniversary, ZT Systems combines individualized attention to your requirements with the experience and scalability you demand.



SC1108i-41

ZT Reliant SC1108i Small Business Server

Supports Intel® AMT 5.0 Professional, Intel® Trusted Platform Module, and is Microsoft Windows Vista Premium WHQL certified.

- Intel® Core™ i7 Processor 920
- 3GB DDR3 1333MHz SDRAM (3x1GB)
- 1TB SATAII 32MB Cache Hard Drive (Max.4)
- 16x DVD-ROM
- nVidia GeForce 9400GT DDR2 DVI/VGA/HDMI
- RedHat Enterprise 5 Trial Version
- 3 Year Limited Warranty and 24x7 Lifetime Toll-Free Phone Support

Starting at only

\$899^{99*}



1221Ri-84

ZT Systems 1221Ri 1U Rack Server

Designed for entry-level file/web/mail servers and small business applications

- Single Intel® Xeon® Processor X5405 (dual capable)
- 4GB DDR2 ECC/REG. SDRAM (2x2GB, Max. 48GB)
- (2) 500GB SATAII 32MB Cache Hard Drives (RAID 0,1,5,10 supported)
- 16x DVD-ROM
- 2 x 3.5" SAS/SATA Hotswap Drive Bays
- RedHat Enterprise 5 Trial Version
- 3 Year Limited Warranty and 24x7 Lifetime Toll-Free Phone Support

Starting at only

\$1099^{99*}

Featuring the **NEW**
Intel® Xeon® Processor 5500 Series



1222Ri-86

ZT Systems 1222Ri 1U Rack Server

Ultimate Performance and Energy Efficiency with new Intel® Xeon® Processor 5500 Series

- Single Intel® Xeon® Processor E5504 (dual capable)
- 6GB DDR3 1333MHz ECC/REG. SDRAM (3x2GB)
- (2) 1TB SATAII 32MB Cache Hard Drives (RAID 0,1,5, 10 supported)
- 16x DVD-ROM
- 4 x 3.5" SAS/SATA Hotswap Drive Bays
- RedHat Enterprise 5 Trial Version
- 3 Year Limited Warranty and 24x7 Lifetime Toll-Free Phone Support

Starting at only

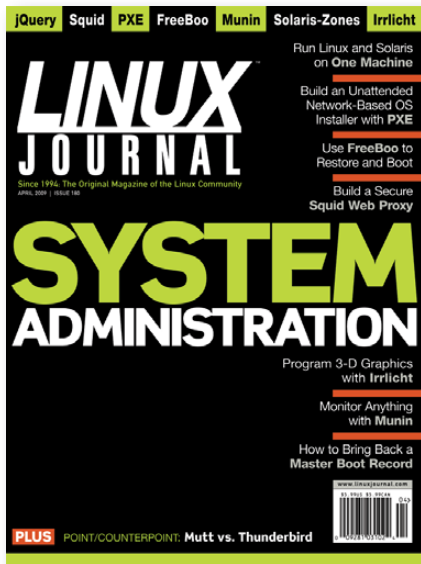
\$1499^{99*}

www.ztsystems.com/linuxjournal

866-984-7687



* Plus shipping and handling. System images may include equipment and/or accessories which are not standard features. Not responsible for errors in typography and/or photography. All rights reserved. All brands and product names, trademarks or registered trademarks are property of their respective companies. Intel, the Intel logo, Xeon, Xeon Inside, Intel Core, and Core Inside are trademarks of Intel Corporation in the U.S. and other countries.



Better Method of Last Resort

There is a better “method of last resort” than the one mentioned in Kyle Rankin’s excellent article in the March 2009 issue titled, “When Disaster Strikes: Hard Drive Crashes”.

If you can’t mount a ddrescue image, but need to retrieve documents, photographs, PDF files and so on, you can use a nifty program called foremost. It is available for most *nix platforms, and on Windoze via cygwin. foremost scans through a hard drive image, mountable or not, and looks for recognizable file headers. It understands more than 20 popular file headers, including jpg, pdf, doc, xml and so on. When it finds these files, it dumps them out as usable files. It is truly a thing of beauty. The first time you use it, you will just sit back in amazement. For example, if you have a hard drive image named my_hd_image.dd that was made with one of the dd utilities, you could execute the following command:

```
~$ foremost -t all -i my_hd_image.dd
```

After the command executes, a subdirectory will be created that has all of the recovered files, organized neatly by file type. On Ubuntu, you can get foremost by typing:

```
~$ sudo apt-get install foremost
```

This tool is also excellent for recovering deleted files from USB drives and more. Enjoy! Love your magazine!

--
mmueller

Wow, great tip. Thanks!—Ed.

Cool Project

I wanted to put in another plug for all the authors of *LJ* and say that you guys and gals do a bang-up job. I recently renewed my subscription, not for one but two years. It is the highlight of my month to find it in the mailbox.

Compliments aside, I have a neat suggestion. I have been scanning forums, journals and search engines for a way to utilize a laptop as a more-or-less standalone DVD player for our vehicle. I know that the logical thing to do would be to use the laptop as a laptop, but I am wanting to get a little more “geeky” and see if I could ceiling-mount something like a 12” unit, link it into the car stereo system through an auxiliary channel and use it as a media server—and even get a wireless mouse/keyboard to “pimp up my ride” a little bit. The thing is, I would like it to be easy to use for my seven- and/or five-year-old and non-techy spouse.

Thanks again to all contributors. I appreciate your knowledge and philosophy.

--
Dean Anderson

Thanks for the compliments. Regarding your car-puter, um, AWESOME! That's a cool project if I've ever heard one. It might be worth investing in touchscreen technology and something like XBMC. Your kids would be able to manage things by tapping. (And, my wife just made it very clear to ME that I'm not allowed to try such a thing with our minivan, so you'll have to keep us posted.)—Ed.

Laconica

Regarding the *LJ* Web article “What’s the Tweeting Protocol?” by Doc Searls

(www.linuxjournal.com/content/whats-tweeting-protocol), is *LJ* ever going to create a social blog for Linux users to come together and work on rallying up new users and clients for Linux as a better and more adaptable OS? Love the mag and its articles.

--
C Anaman

We always are interested in how best to take over the world. Unfortunately, we still haven't figured out the secret sauce for making it take off. Please don't hesitate to drop us suggestions; we're always open to ideas.—Ed.

Misery Loves Company

I work for an agency of the California State government. Our agency is going open source as much as possible. Our programmers even slap the GPL on everything they make at work. We have Linux servers and use MySQL, Apache and nginx, among other open-source applications. There’s really only one major hurdle stopping us from putting Linux on the desktop: specialized applications. For example, there’s an application that allows a judge to click on case characteristics, and the app then spits out everything the judge is supposed to tell the jury before the trial begins. The laws that govern what judges are supposed to tell juries change quarterly; therefore, we receive quarterly updates.

Unfortunately, like many of our other specialized applications, it doesn’t work on Wine. There are no open-source or even commercial Linux programs that do the same thing.

I’ve thought of virtual machines, but that still requires us to purchase Windows licenses, as would dual-booting or terminal services/Citrix.

I am hoping that as Linux advocates and experts, you could suggest a strategy that would allow Linux on the desktop at our organization given our circumstances.

--
Toby Richards

I deal with this every day working at a school district. You're absolutely correct that a few proprietary applications are tough to deal with. So far, the best way I've managed to handle it is with a single Windows terminal server accessed remotely via rdesktop. You don't need to pay for the CAL for the Linux machine, but you still do need to pay for the server and TSCALs. I think, ultimately, the best hope will be as programmers continue to move their applications to Web-based alternatives. I think that will truly level the playing field on the desktop.

Thanks again for the question, and although I'm not much help, perhaps misery loves company.—Ed.

My Carbon Footprint Is Doing Just Fine, Thank You

In James Gray's "Go Green, Save Green with Linux" article in the April 2008 issue, and again in a response to that article in the April 2009 Letters section titled "Ouch!", we continue to be misled into believing that carbon dioxide is a deadly poison. This is irresponsible journalism, or journalism without research to back up the statements. Carbon dioxide is, in fact, an essential ingredient to life on this planet. It produces oxygen that we breathe through a process called photosynthesis, and without carbon dioxide, we would suffocate. Additionally, the whole "carbon footprint" scam is simply a fraud designed by fear-mongers whose aim is to introduce a "carbon tax" against your "carbon footprint"; it's all the same scam as the global warming myth, and it's simply political fear mongering. If you research it, the carbon dioxide output from humans is miniscule compared to the carbon dioxide output from the earth itself. For us to believe that we are having an impact on global temperatures via our carbon dioxide output is absolutely absurd. Please, LJ, let's all stop perpetuating these myths. Don't get me wrong; I'm all for conservation of energy and all the benefits from that, but let's get our facts straight. Thanks.

--
Mike

James Gray replies: Thanks for writing. You are completely correct that carbon dioxide is essential to life on Earth. For the record, nowhere did I say that it is a deadly poison. However, I do contend that an imbalance of carbon dioxide in the atmosphere appears to have an

effect on our global climate patterns. On one level, it is simple physics. Carbon dioxide is one of many greenhouse gases that trap heat in the lower atmosphere, enabling life to exist. Logically, if more carbon dioxide exists in the atmosphere—and we are filling

ServerBeach

by geeks, for geeks™



When **YouTube** first started to experience its exponential growth and our hosting needs changed, ServerBeach offered us great flexibility. They continually redesigned our streaming architecture for optimum performance while keeping our hosting costs in check.

STEVE CHEN Founder | **YouTube**

”

ValuePack (always included)

- > 24/7 live customer service
- > 24/7 ticketing system
- > Personal account manager
- > Lots of bandwidth
- > Free OS reloads
- > Free Rapid Reboot
- > Free Rapid Rescue
- > Super fast PEER 1 network
- > Rock-solid IT infrastructure
- > 100% uptime guarantee
- > Choose your data center - East Coast, West Coast and Central

serverbeach.com **1.800.741.9939**

A PEER 1 COMPANY

[LETTERS]

it with around 30 gigatons of a gas every year—more heat will be trapped. However, this is not something I pulled out of a hat. It is called the Theory of Global Climate Change. To learn about the theory, I have read documents, such as the “Climate Change 2007: Synthesis Report” from the Intergovernmental Panel on Climate Change (IPCC). The IPCC, which recently won the Nobel Peace Prize for its work, summarizes the scientific findings of climatologists around the world. This report states: “Warming of the climate system is unequivocal, as is now evident from observations of increases in global air and ocean temperatures, widespread melting of snow, ice and rising average sea level.” Regarding causes, the report says: “Global GHG [greenhouse gas] emissions due to human activities have grown since pre-industrial times, with an increase of 70% between 1970 and 2004....Most of the observed increase in global average temperatures since the mid-20th century is very likely due to the observed increase in anthropogenic [that is, human-caused] GHG

concentrations.” Read the full document at www.ipcc.ch/pdf/assessment-report/ar4/syr/ar4_syr_spm.pdf.

I’m curious how you’ve determined that a scientific theory—one supported by huge amounts of empirical data and having near unanimous consensus among climatologists—is a “myth”. Are you a climatologist who has collected his own data? Do you reject other scientific theories, such as plate tectonics, the Big Bang or relativity? Or do you just reject those that are inconvenient to you?

I periodically write about the Theory of Global Climate Change because I care about the planet I will leave to my descendants, and I am fascinated by the wonders of nature. I also will stridently advocate for a carbon tax, because such tools are the only effective way to change human behavior. I think that if you do your homework, you will find yourself on the same team—the one that is informed by today’s best science and advocates for prevention today in order to avoid a future climate catastrophe.

Poor Little Nerd Boy

I have been watching your videos on YouTube for quite some time now, and I have become so interested in Linux, I have installed it on my PS3. (I know you’re thinking, how can you be poor if you own a PS3? Well, the answer for that is my friend sold me his for \$30, because he got a 360.) Anyway, I have sent you an e-mail because I really would like to get a fully functional Ubuntu Linux computer of some sort. I was wondering if you could send me one of your test laptops that you receive? My parents tell me that I should get a job to earn the money for one, but sadly, I am not old enough for a job, so I became desperate and have turned to you. I don’t expect to be e-mailed back, but it would make my day if you would reply.

--
Alexander

Sadly, we don’t get to keep most of the products we review. That said, if a fully functional Linux machine is really something you want, it might be worth checking with the local schools or government buildings. Often, last year’s models are either given away or sold for very good prices. If you draft a convincing letter to the right administrator at a local school district, you might be surprised with the response you get. Good luck!—Ed.

Keep On Keeping On

Today I got a letter saying that *Dr. Dobbs*, the other magazine I’m subscribed to, ceased as a standalone monthly magazine. I hope you’re not planning to do the same. *LJ* is the main/only source of information I have on Linux, and I certainly would miss not receiving it anymore. Keep up the good work!

--
Pedro

We certainly don’t plan to go anywhere! If you now have more in your subscription budget, perhaps you could get a subscription to Linux Journal for someone else. (Shameless plug?) Also, don’t forget to stop by our Web site. We have tons of additional content there as well.—Ed.

Live CDs

In the Linux world, many things have been changing during the past few years. The boom that had started with

PHOTO OF THE MONTH

Have a photo you’d like to share with *LJ* readers? Send your submission to publisher@linuxjournal.com. If we run yours in the magazine, we’ll send you a free T-shirt.



This was just a few of all the plush penguinistas at the Penguin exhibit shop at SeaWorld in San Diego. At least there is not a Window or Apple amongst them!
Submitted by Curtis Vaughan.

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 +1-818-487-2089. E-mail us at subs@linuxjournal.com or reach us via postal mail, Linux Journal, PO Box 16476, North Hollywood, CA 91615-9911 USA. Please remember to include your complete name and address when contacting us.

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

LETTERS TO THE EDITOR: We welcome your letters and encourage you to submit them at www.linuxjournal.com/contact or mail them to Linux Journal, PO Box 980985, Houston, TX 77098 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 +1 713-344-1956 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.

Tip

I'm right behind Fritz Mehner's desire [see the March 2009 Letters section] to make scripting more elegant and efficient wherever possible.

Fritz's letter drew my attention to a problem with the example for loop at the end of the Tech Tip referred to by him (*LJ*, December 2008, page 56), which needs the opposite treatment if scripted to avoid potential disaster. The problem is the "globbed" for statement, which will fail if the total length of matching filenames causes the expanded command line to exceed the shell's command-line buffer. Although that buffer is very large, it is finite (of the order of 10,000 bytes usually), and I've fixed more than one production script that had failed due to this defect.

The solution is to replace the first line of the loop:

```
for file in *; do
  data_source $file | ...
done
```

with one that uses the find utility to avoid shell expansion, as in the replacement loop below:

```
find . -maxdepth 1 -type f -print | while read file; do
  data_source $file | ...
done
```

I hope this helps someone avoid that otherwise inevitable call in the middle of the night.

--
Ross Johnson

Knoppix, a live CD that was easy to use and made it possible to learn about Linux and use it without being an expert, changed everything. Now there are hundreds of distributions, many of which can be used as live CDs. But, it became difficult to have a clear overview of all of them. Many are good, and some are specialized, such as GoboLinux (KDE, de/engl/mag/port). In GoboLinux, you don't need a package database, because the filesystem is the database; each program resides in its own directory. There's also OpenGEU (Enlightenment, Italian/English), openmamba (KDE, Italian/English, good for beginners), Puppy (JWM, manual, English, small, quick, easy to use, good on old and new hardware), SAM (Xfce, multilingual), SliTaz (JWM, English/French, extra small) and so on.

It's fascinating what can be offered by mini-distributions—they're fast, easy to install and extremely quick. DSL is well known in the Linux world, but there are more, such as Puppy, Feather, Slax and, last but

not least, SliTaz, the smallest distribution I know of at the moment that includes programs like an Internet browser.

I think it would be a great idea to publish the strength of live CDs and show the developments, especially with the very small distributions that do a great job, and demonstrate that there are not only alternatives to Windows but also alternatives to the big ones like Sabayon, Mandriva, Ubuntu and SUSE.

--
computerophil

Unfortunately, as Linux users, one of our biggest strengths is also one of our biggest weaknesses. As you've outlined, "Linux" for the end user doesn't come close to describing the desktop experience. We do try to cover a wide variety of distributions, but there are so many options, it's hard to cover everything equally. Be sure to check out our Web site as well (www.linuxjournal.com). It's another way to spread the attention around a bit.—Ed.

diff -u

WHAT'S NEW IN KERNEL DEVELOPMENT

An effort to change the license on a piece of code hit a wall recently. **Mathieu Desnoyers** wanted to migrate from the **GPL** to the **LGPL** on some userspace **RCU** code. Read-Copy Update is a way for the kernel to define the elements of a data object, without other running code seeing the object in the process of formation. Mathieu's userspace version provides the same service for user programs. Unfortunately, even aside from the usual issue of needing permission from all contributors to change the license of their contribution, it turns out that **IBM** owns the patent to some of the RCU code concepts, and it has licensed the patent for use only in GPLed software. So, without permission from IBM, Mathieu can get permission from all the contributors he wants and still be stuck with the GPL.

Loadlin is back in active development! The venerable tool boots Linux from a directory tree in a **DOS** partition, so all of us DOS users can experiment with this new-fangled Linux thing. To help us with that, **Samuel Thibault** has released Loadlin version 1.6d and has taken over from **Hans Lerman** as official maintainer of the code. The new version works with the latest Linux kernels and can load up to a 200MB bzimage. He's also migrated

development into a **mercurial** repository. (Although not as popular as git with kernel developers, mercurial does seem to have a loyal following, and there's even a book available at **hgbook.red-bean.com**.) After seven years of sleep, here's hoping Loadlin has a glorious new youth, with lots of new features and fun. It loads Linux from DOS! How cool is that?

Hirofumi Ogawa has written a driver for **Microsoft's exFAT filesystem**, for use with large removable Flash drives. The driver is read-only, based on reverse-engineering the filesystem on disk. There doesn't seem to be immediate plans to add write support, but that could change in a twinkling, if a developer with one of those drives takes an interest in the project. Hirofumi has said he may not have time to continue work on the driver himself.

Meanwhile, **Boaz Harross** has updated the **exofs filesystem**. exofs supports **Object Storage Devices (OSDs)**, a type of drive that implements normal block device semantics, while at the same time providing access to data in the form of objects defined within other objects. This higher-level view of data makes it easier to implement fine-grained data management and security. Boaz's updates include some ext2 fixes that still apply to the


exofs codebase, as exofs originally was an ext2 fork. He also abandoned the **IBM API** in favor of supporting the **open-osd API** instead.

Adrian McMenamin has posted a driver for the **VMUFAT filesystem**, the **SEGA Dreamcast** filesystem running on the Dreamcast visual memory unit. Using his driver, he was able to manage data directly on the Dreamcast. At the moment, the driver code does seem to have some bugs, and other problems were pointed out by various people. Adrian has been inspired to do a more intense rewrite of the code, which he intends to submit a bit later than he'd first anticipated.

A new source of controversy has emerged in Linux kernel development. With the advent of pocket devices that are intended to power down when not in use, or at least go into some kind of power-saving state, the whole idea of suspending to disk and suspending to RAM has become more complicated. It's not obvious whether the kernel or userspace should be concerned with analyzing the **sleep-worthiness** of the various parts of the system, or how much the responsibility should be shared between them. There seems to be many opinions, all of which rest on everyone's idea of what is appropriate as well as on what is feasible. The kernel

is supposed to control all hardware, but the X Window System controls hardware and is not part of the kernel. So, clearly, exceptions exist to any general principles that might be involved. Ultimately, if no obvious delineation of responsibility emerges, it's possible folks may start working on competing ideas, like what happened initially with software suspend itself.

—ZACK BROWN



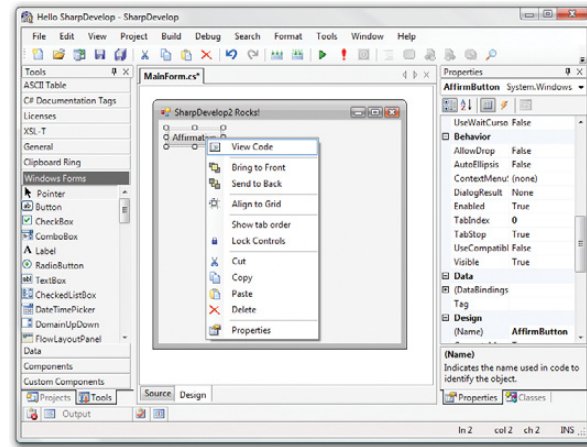
LJ Index June 2009

- Number of projects listed on sourceforge.net: **161,671**
- Number of new projects registered on sourceforge.net during the first four days of March 2009: **278**
- Number of projects updated on sourceforge.net during the first four days of March 2009: **755**
- Number of projects on sourceforge.net found by searching for "Linux": **10,222**
- Number of projects on sourceforge.net found by searching for "Windows": **7,813**
- Number of projects on sourceforge.net found by searching for "Mac OS X": **5,751**
- Percentage of time you type sourceforge.net rather than sourcefor.net the first time: **98.2%**
- Number of projects listed on ohloh.net: **275,412**
- Number of people listed on ohloh.net: **245,752**
- Start date (in UNIX time) of GCC repository, the third oldest tracked by ohloh.net (November 1988): **594,367,200**
- Start date (in UNIX time) of GNU Emacs repository, the second oldest tracked by ohloh.net (April 1985): **481,183,200**
- Start date (in UNIX time) of BRL-CAD, the oldest tracked by ohloh.net (April 1983): **418,024,800**
- Number of projects listed on freshmeat.net: **45,831**
- Number of users listed on freshmeat.net: **412,744**
- Number of projects listed on code.google.com: **50,000+**
- Number of projects listed on Microsoft's open-source project site codeplex.com: **8,030**
- Number of projects updated on codeplex.com during the first four days of March 2009: **101**
- Number of Google hits for "open source code repositories": **47,100,000**
- US National Debt as of 03/05/09, 9:36:45am CST: **\$10,950,269,741,924.01**
- Change in the debt since last month's LJ Index: **\$174,023,143,132.25**

Sources: 1-6: sourceforge.net | 7: *Thin air*
8-12: ohloh.net | 13, 14: freshmeat.net | 15: code.google.com
16, 17: www.codeplex.com | 18: www.google.com
19: www.brillig.com/debt_clock | 20: *Math*

NON-LINUX FOSS

In our second Upfront installment highlighting non-Linux FOSS projects, we present SharpDevelop. SharpDevelop (aka #Develop) is an IDE for developing .NET applications in C#, F#, VB.NET, Boo and IronPython. SharpDevelop includes all the stuff you'd expect in a modern IDE: syntax highlighting, refactoring,



SharpDevelop Running on Vista (from www.icsharpcode.net)

At the time of this writing, version 3.0 just recently has been released.

Even if you use only Linux, you may be indirectly using SharpDevelop. If you use any Mono programs, they probably were developed using the MonoDevelop IDE. MonoDevelop was forked from SharpDevelop in 2003 and ported to GTK.

—MITCH FRAZIER

Delivering Content to Your Desktop with Miro

I've been a fan of Miro since it originally came out as a program titled Democracy. This open-source, cross-platform project has evolved over the years into an almost perfect example of how to watch on-line media. Miro can play almost any non-DRM video format. What makes it really unique, however, is its ability to retrieve fresh content automatically. It supports the traditional on-line media providers, like Revision 3, but it also adds the ability to subscribe to any RSS feed of videos or even torrents of videos. Miro will download torrent files automatically with its built-in BitTorrent client from any RSS feed you throw at it. Add sites like Hulu.com to Miro's arsenal of content, and you have an almost perfect video-watching experience.

Keep watching this project, because although there isn't yet an interface that's easy to control from a couch, that's a pretty simple change that would make Miro a candidate for your living-room television portal. As it is right now, it's a great addition to anyone's computer desktop. In fact, you can subscribe to *Linux Journal's* videos and have them delivered to your desktop automatically. I'm not sure if being stuck with my face on your desktop every week is a very good selling point, but at least there isn't a subscription fee. Everything Miro offers, including its software, is completely free. Check it out at www.miro.org. —SHAWN POWERS

They Said It

Linux is a cancer that attaches itself in an intellectual property sense to everything it touches.

—Steve Ballmer, June 2001

Will we interoperate with products that come, like Linux, from the Open Source world? Yes, we will. Will we encourage people who want to do open-source development to do it on top of Windows? Yes.

—Steve Ballmer, July 2008

Sun's doing tremendous damage to the project.

—Geir Magnusson Jr, referring to a Java licensing dispute between Sun and the Apache Software Foundation

Religion-themed domains could provoke "bitter disputes" that would force ICANN into "recognizing to a particular group or to a specific organization the legitimacy to represent a given religious tradition".

—Monsignor Carlo Maria Polvani, in a letter to outgoing ICANN chief Paul Twomey, concerning new Internet domains, such as .catholic, .islam, .muslim and so on, or as *The Register* called them, the dot god domains.

Just because something doesn't do what you planned it to do, doesn't mean it's useless.

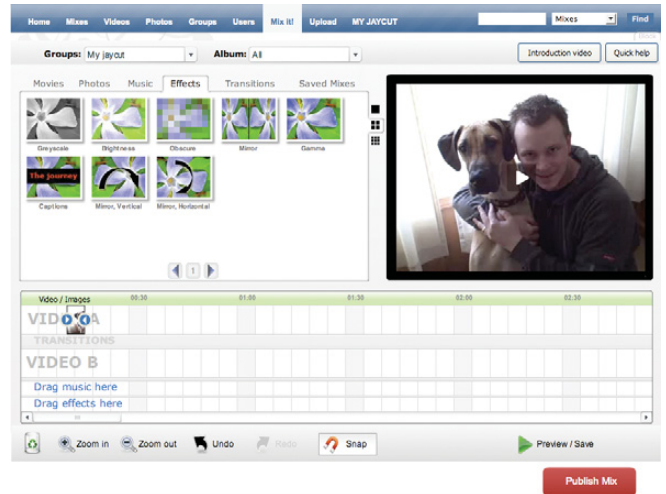
—Thomas Alva Edison

THE WEB, MAKING YOUR COMPUTER OBSOLETE

As a video creator, I often try as many different video editing options I can find. I've used video editors on every platform, in every style and every design. I never thought such a resource-intensive process would be able to move to the Web, but as is all too often the case, I was wrong.

If you are like me, you never even considered searching for on-line alternatives for video editing software. You might want to reconsider. I haven't had a chance to test them all, but a simple Google search for "on-line video editor" provides a ton of options. Most of them are completely free, and a few are rather robust. Jaycut (www.jaycut.com), for example, has a look that almost rivals desktop video editing software. Although on-line video editors haven't quite surpassed the abilities of their desktop counterparts, I was more than impressed by how far they've come. Because on-line video editing means you have your tools with you wherever you go, the future of video editing might look very different. It sure beats toting around an external hard drive full of raw DV footage.

—SHAWN POWERS



Jaycut offers some great editing features right inside your browser.

LinuxJournal.com

As you read through this year's Readers' Choice Award winners, I'm sure you'll find a few items you'll want to learn more about. You'll discover a wealth of information about almost anything on the list at LinuxJournal.com. With more than 15 years of articles, you'll find what you're looking for, and maybe even learn how some of these tools have progressed over the years.



Head to LinuxJournal.com, and take the search box for a spin. You'll see all your favorites like Python, Nagios and Ubuntu. To get started, check out your favorite game, *Frozen Bubble*. Shawn Powers gives you a little preview:

www.linuxjournal.com/video/better-solitaire-frozen-bubble.

—KATHERINE DRUCKMAN

Linux in a Minute

Are you new to Linux? Are you an old hand, but want to brush up on your command-line skills? Heck, are you just someone who likes to learn about the tips we dream up here at *Linux Journal*? Our new video series on LinuxJournal.com brings you new video tips almost every day. Here are some of the topics we've recently covered:

- Unetbootin (Bootable USB Linux Installer): www.linuxjournal.com/video/creating-bootable-usb-install-drives-unetbootin
- Command-Line 101 (grep, top, ls): www.linuxjournal.com/video/commandline-101-getting-grip-grep, www.linuxjournal.com/video/commandline-101-using-top and www.linuxjournal.com/video/commandline-101-basic-directory-commands
- Extract MP3 from a Video: www.linuxjournal.com/video/extract-mp3-audio-portion-video
- Donating CPU Cycles with BOINC: www.linuxjournal.com/video/donate-cpu-cycles-boinc
- Installing VirtualBox: www.linuxjournal.com/video/installing-linux-virtual-box



- Using the screen Command: www.linuxjournal.com/video/transfer-your-terminal-screen

With videos coming out almost every day, you're bound to find something of use in our short one-minute tutorials. Check them out at www.linuxjournal.com/linux-minute.

—SHAWN POWERS

SMALL, EFFICIENT COMPUTERS WITH PRE-INSTALLED UBUNTU.

GS-Lo8 Fanless Pico-ITX System

Ultra-Compact, Full-Featured Computer
Excellent for Industrial Applications



3677 Intel Core 2 Duo Mobile System

Range of Intel-Based Mainboards Available
Excellent for Mobile & Desktop Computing



DISCOVER THE ADVANTAGE OF MINI-ITX.

Selecting a complete, dedicated platform from us is simple: Pre-configured systems perfect for both business & desktop use, Linux development services, and a wealth of online resources.



LOGIC
SUPPLY

www.logicsupply.com



REUVEN M. LERNER

Checking Your HTML

Integrate HTML validation into your test suite for better HTML from the get-go.

We say that Tim Berners-Lee invented the World Wide Web, and that's certainly true. But, we can boil the Web down to three specific technologies: URLs (for uniquely identifying resources on the Internet), HTTP (a stateless protocol for transmitting documents) and HTML (a markup language). Each of these inventions was simple to understand, as well as simple to implement. And, it is this combination of simplicity and elegance that has made the Web the success that it is.

All three of these technologies have evolved over the years, reflecting new uses and needs. For example, HTTP now supports a system of "headers" in both the request and response, which can do everything from indicate the content type of the response body to provide hints regarding how long the data should be cached.

HTML has grown up quite a bit as well, evolving to become a truly semantic markup language (with styling information moved to external CSS documents) with a more rigorous and standardized definition. Standardization has made HTML slightly harder to write, in that you need to be more careful about items, such as tag names (keeping them lowercase), attributes (because not all are valid in all contexts) and closing tags. One advantage to such standardization is that we now can predict to a much greater degree what pages will look like across different browsers. Sloppy HTML means that the browser has to decide what you meant, which can have consequences that vary widely in their influence on the way the page looks.

More significant, the rise of AJAX as a paradigm for Web development has made it increasingly important that HTML be well formed. Many AJAX-related routines need to modify a particular element on the page in some way. The easiest way to do this typically is to grab the element via its id attribute, which is guaranteed to be unique. (If you want more than one element to use an ID, you really should use a class instead.) In the last few months, I have worked on a number of pages that had duplicate ID attributes. Sometimes this was the result of a simple mistake, and sometimes it resulted from ignorance on the part of a Web designer. But in all cases, this meant that my JavaScript performed differently from what I expected.

Although HTML validation might seem boring, it's actually an essential part of getting AJAX-powered,

latest-paradigm, super-fancy Web sites to work. This month, I review a few tools I use to make sure the HTML I create is as standards-compliant as usual. I begin with some simple, manual tests that can run on individual pages. Then, I show some automated tools I use when developing applications in Ruby on Rails, allowing me to check the HTML of all of my pages en masse, including those that require password protection to access.

HTML Standardization

Before continuing, it's important to realize that HTML is a catchall term for many different, related markup languages. And, when I say markup, I mean that HTML is a language used to describe a text, identifying its different parts. For instance, a newspaper article will have a headline, one or more authors, one or more paragraphs of text, zero or more photographs, and one or more captions per photograph. A markup language doesn't add content to a document, but rather describes the individual parts of the document, so that they can be laid out and displayed in an appropriate way. In this sense, HTML is a direct descendant of SGML, a markup language that was developed many years previously, but which was far more difficult to work with.

Although there have been several versions of HTML over the years, let's focus on the ones that are most widely used today. Perhaps the most common version of HTML is an unstructured, unversioned, nonstandard document. I'm certainly guilty of creating many such documents, which look like this:

```
<html>
<head>
  <title>This is the title</title>
</head>
<body>
  <h1>This is the headline</h1>
  <p>This is a paragraph</p>
  <p>This is another paragraph</p>
</body>
</html>
```

Nothing is wrong, per se, with the above document. But, because it fails to indicate which version of HTML it is using, browsers must make a variety of assumptions. These assumptions can make it

hard to predict how different browsers will operate, using something known as quirks mode.

Fortunately, we can choose a standard implementation and indicate that to a browser by adding a DOCTYPE declaration at the top of the document. When assigning the value of DOCTYPE, you need to decide whether you will use HTML or XHTML (that is, an XML-compliant version of HTML), and whether you want the strict, transitional or frameset variety of that markup language.

The strict version of each markup language is the ideal version that allows no styling elements. On a modern site, such styling should be defined in CSS, not in HTML. However, it may be difficult for some sites to comply with the strict definition, either because their authoring tools use tags that aren't allowed in the strict definition, or because the site's authors want to use forbidden elements, such as those for embedded Flash. To make the transition to strict HTML easier, the standards allow for transitional HTML, which provides a larger number of tags.

Let's define our tiny document as follows:

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01 Transitional//EN">
<html>
<head>
  <title>This is the title</title>
</head>
<body>
  <h1>This is the headline</h1>
  <p>This is a paragraph</p>
  <p>This is another paragraph</p>
</body>
</html>
```

The `<!DOCTYPE>` declaration at the top of the page tells browsers (and any other programs that might try to parse the page) that we want to follow the standards, but that we'll do so using the transitional declaration.

Once we have indicated our willingness to apply the transitional standard, we may discover that our documents are no longer valid. For example, if I include an image in my HTML document:

```

```

Powerful: Rhino



Rhino M6400/E6500

- Dell Precision M6400/ Latitude E6500
- 2.2-3.0 GHz Core 2 Duo or 2.5 GHz Core 2 Quad
- Up to 17" WUXGA LCD w/ X@1920x1200
- NVidia Quadro FX 3700M
- 80-500 GB hard drive
- Up to 16 GB RAM
- DVD±RW or Blu-ray
- 802.11a/g/n
- Starts at \$1360

- High performance NVidia 3-D on a WUXGA widescreen
- High performance Core 2 Quad, 16 GB RAM
- Ultimate configurability — choose your laptop's features
- One year Linux tech support — phone and email
- Three year manufacturer's on-site warranty
- Choice of pre-installed Linux distribution:



Tablet: Raven



Raven X200 Tablet

- ThinkPad X200 tablet by Lenovo
- 12.1" WXGA w/ X@1280x800
- 1.2-1.86 GHz Core 2 Duo
- Up to 8 GB RAM
- 80-320 GB hard drive / 128 GB SSD
- Pen/stylus input to screen
- Dynamic screen rotation
- Starts at \$2410

Rugged: Tarantula



Tarantula CF-30

- Panasonic Toughbook CF-30
- Fully rugged MIL-SPEC-810F tested: drops, dust, moisture & more
- 13.3" XGA TouchScreen
- 1.6 GHz Core 2 Duo
- Up to 8 GB RAM
- 80-320 GB hard drive
- Call for quote

Save up to \$200. Details at:
www.EmperorLinux.com/lj0609

EmperorLinux

...where Linux & laptops converge

www.EmperorLinux.com

1-888-651-6686



Model specifications and availability may vary.

With the above line inserted into my document, it is no longer valid, because it is missing an alt attribute. Once I add that attribute, the document is valid:

```

```

However, we can get even better results if we enforce XML considerations and declare our document to be XHTML transitional. To do that, we modify not only the !DOCTYPE declaration, but also the <html> tag:

```
<!DOCTYPE html
    PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">

<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en" lang="en">
  <head>
    <title>This is the title</title>
  </head>
  <body>
    <h1>This is the headline</h1>
    <p>This is a paragraph</p>
    <p>This is another paragraph</p>
    
  </body>
</html>
```

Suddenly, our document is invalid again. Because we have declared it to be XHTML transitional, we need to follow XML rules. We need to close our tag, most easily accomplished by using the self-closing syntax:

```
<!DOCTYPE html
    PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">

<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en" lang="en">
  <head>
    <title>This is the title</title>
  </head>
  <body>
    <h1>This is the headline</h1>
    <p>This is a paragraph</p>
    <p>This is another paragraph</p>
    
  </body>
</html>
```

With that in place, our document is now valid. As you can imagine, finding all the problems that might occur in a document can be difficult, even for someone who is trained and experienced. Trying to check all the pages on a site, particularly one that contains hundreds or thousands of pages, clearly would be impossible.

The solution, then, is to have a program check the pages' validity automatically, preferably as part of your automated tests. This way, you can discover when you have problems quickly and easily.

W3C Validator

One of the best tools for checking the validity of a page's markup is the World Wide Web Consortium's validator, available at validator.w3.org. I use the validator almost exclusively from within Firefox, into which I have installed the Web Developer plugin. This plugin lets you validate the HTML of any page, simply by selecting Validate HTML from the browser. The browser submits the page's URL to the W3C validator, which then gives a line-by-line indication of what problems (if any) the page contains.

The W3C validator has at least two problems, however. First, it requires that you submit each page, one at a time, to the validator program. This means a great deal of time and effort, just to check your pages. A second consideration is more practical; the validator works only with pages that are accessible via the Internet, without password protection. If your site is being developed on your local computer, and if you have a firewall protecting your business from the outside world, you probably will be unable to use the validator via the Web.

One solution to this problem is to install the W3C validator on your local computer. You can get the source code from validator.w3.org/source, which comes in the form of a Perl program. On modern Debian and Ubuntu machines, you can install `w3c-markup-validator`, which makes it available via your local Web server, ready to be invoked.

If you end up installing the validator manually, it requires a number of modules, which you might need to download from CPAN (Comprehensive Perl Archive Network), a large number of mirrors containing open-source Perl modules. It might take some trial and error to figure out which modules are necessary, although if you are an experienced user of the CPAN.pm installer, this shouldn't be too much trouble. Note that the `SGML::Parser::OpenSP` module requires the OpenSP parser, which you can get from SourceForge at openjade.sf.net.

As you might be able to tell, a number of these modules are required in order to handle alternate encoding schemes, particularly those for Asian languages. Even if you aren't planning to handle such languages, the modules are mandatory and must be installed.

The validator program, called `check`, should be put in a directory for CGI programs or in a directory handled by `mod_perl`, the Apache plugin that lets you run Perl programs at a higher speed, among other things. You also will need to install a configuration file, typically placed in the directory `/etc/w3c`,

but which you can relocate by setting the W3C_VALIDATOR_CFG environment variable.

Validating Rails Templates

Now that you have the W3C checker installed on your own server, you can feed it URLs that aren't open to the public. But, if you are developing an application in Ruby on Rails, you can go one step better than this, integrating the W3C validator into your automated testing.

In order to do this, you need to install the html_test plugin for Rails. Go into your Rails application's root directory, and type:

```
script/plugin install
➔http://htmltest.googlecode.com/svn/trunk/html_test
```

With this plugin in place, you now can use three new assertions in your functional and integration tests: `assert_w3c` returns true if the W3C validator approves of your HTML; `assert_tidy` returns true if you're using the HTML Tidy library, described below; and, `assert_validates` calls both of these.

So, if you have a FAQ page you want to check with an integration test, you can write something like this:

```
def test_faq
  get '/faq'
  assert_response :success
  assert_w3c
end
```

If the HTML for this page is approved by the W3C validator, everything is fine. If this page is not valid, you will get quite a bit of output, which you should redirect to a file. This file will contain not only the results of your tests, but also the same HTML output that you would have gotten from the public, Web-based W3C validator. This means you'll get a complete and easy-to-read description of what you did wrong.

You'll often discover that a large number of validation errors can be fixed with a small number of corrections. For example, when I ran this test against a sloppy FAQ page, I got six validation errors. I was able to fix all of them by indicating the appropriate namespace in my `<html>` tag and removing an extraneous `</p>` from the end of the file.

Checking HTML validity in this way is nice and easy. (It can be time consuming, however, to invoke the validator on every single page; I think the trade-off is worthwhile, but you might disagree.) If you always want to check HTML validity, you can change your test environment's configuration somewhat, so that it'll happen automatically, without having to invoke `assert_w3c` each time.

To do this, you need to modify `test_helper.rb`, which sits at the top of the test directory, and which is included into every test program. All you have to do is add:

```
ApplicationController.validate_all = true
ApplicationController.validators = [:w3c]
```



Want your business to be more productive?

The ASA Servers powered by the Intel Xeon Processor provide 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.)

1U Server - ASA1401i

- 1TB Storage Installed. Max - 3TB.
- Intel Dual core 5030 CPU (Qty-1), Max-2 CPUs
- 1GB 667MGZ FBDIMMs Installed.
- Supports 16GB FBDIMM.
- 4X250GB htswap SATA-II Drives Installed.
- 4 port SATA-II RAID controller.
- 2X10/100/1000 LAN onboard.



2U Server - ASA2121i

- 4TB Storage Installed. Max - 12TB.
- Intel Dual core 5050 CPU.
- 1GB 667MGZ FBDIMMs Installed.
- Supports 16GB FBDIMM.
- 16 port SATA-II RAID controller.
- 16X250GB htswap SATA-II Drives Installed.
- 2X10/100/1000 LAN onboard.
- 800w Red PS.



3U Server - ASA3161i

- 4TB Storage Installed. Max - 12TB.
- Intel Dual core 5050 CPU.
- 1GB 667MGZ FBDIMMs Installed.
- Supports 16GB FBDIMM.
- 16 port SATA-II RAID controller.
- 16X250GB htswap SATA-II Drives Installed.
- 2X10/100/1000 LAN onboard.
- 800w Red PS.



5U Server - ASA5241i

- 6TB Storage Installed. Max - 18TB.
- Intel Dual core 5050 CPU.
- 4GB 667MGZ FBDIMMs Installed.
- Supports 16GB FBDIMM.
- 24X250GB htswap SATA-II Drives Installed.
- 24 port SATA-II RAID. CARD/BBU.
- 2X10/100/1000 LAN onboard.
- 930w Red PS.



8U Server - ASA8421i

- 10TB Storage Installed. Max - 30TB.
- Intel Dual core 5050 CPU.
- Quantity 42 Installed.
- 1GB 667MGZ FBDIMMs.
- Supports 32GB FBDIMM.
- 40X250GB htswap SATA-II Drives Installed.
- 2X12 Port SATA-II Multilane RAID controller.
- 1X16 Port SATA-II Multilane RAID controller.
- 2X10/100/1000 LAN onboard.
- 1300 W Red Ps.



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 typographic errors.



You also can check the validity of URLs and redirects; although these aren't checking HTML validity per se, they do come with the `html_test` plugin and are quite useful:

```
ApplicationController.check_urls = true
ApplicationController.check_redirects = true
```

With these four lines in your `test_helper.rb`, you can run your integration tests once again. If any of the validation tests fail, you can look at `/tmp/w3c_last_response.html`, which will contain the complete output of that failure. This doesn't help very much if you have multiple failures, however.

If you have designed your templates using the DRY (don't repeat yourself) principle, fixing HTML markup problems shouldn't be too bad. In many cases, you will need to change only one tag in the layout to fix everything.

HTML Tidy

The W3C validator is excellent, but it doesn't always catch everything, such as empty tags. For this, you might want to add to your arsenal, integrating the open-source Tidy library, which identifies and fixes badly written HTML. Tidy originally was written by Dave Raggett, one of the best-known developers

from the early days of the Web; the project is now on SourceForge at tidy.sf.net.

To integrate Tidy checking into your Rails application, first install the library from SourceForge. Then, install the Ruby gem for Tidy integration:

```
sudo gem install tidy
```

Finally, download and install the Rails Tidy plugin:

```
cd vendor/plugins
wget
http://www.cosinux.org/~dam/projects/rails-tidy/rails_tidy-0.3.tar.bz2
tar -jxvf rails_tidy-0.3.tar.bz2
```

Now, modify `test_helper` to read:

```
ApplicationController.validators = [:w3c, :tidy]
```

With that in place, every request to your server now will be checked by both validators, rather than just one.

The Rails Tidy plugin can be useful beyond checking and validating to fix your HTML as it is sent from your server to the user's browser. Although I like this idea in theory, it seems fairly inefficient and slow to parse and rewrite every bit of HTML as it is sent. Plus, I feel that debugging Web applications (and CSS) is tough enough without having the HTML magically rewritten behind the scenes.

Conclusion

HTML has evolved quite a bit over the years, and getting your pages to contain valid HTML can be difficult to handle manually. For this reason, using automated checks and integrating those checks into a Web application's automated settings is a good way to ensure that your site is adhering to HTML standards as closely as possible. This not only gives you the greatest chance of having the site render similarly on different platforms, but it also even may boost your ranking in Google (an assertion I have seen mentioned in several places, but for which I obviously have no proof).

If you are using Ruby on Rails, you can validate your HTML easily from the start of your project. By doing so, you will make life easier for yourself down the line. Moreover, this is far easier than checking pages manually, and it ensures that even administrative and other hidden pages are validated. ■

Reuven M. Lerner, a longtime Web/database developer and consultant, is a PhD candidate in learning sciences at Northwestern University, studying on-line learning communities. He recently returned (with his wife and three children) to their home in Modi'in, Israel, after four years in the Chicago area.

Resources

The Web Developer plugin for Firefox, which has been an invaluable help in my work for several years, is at chrispederick.com/work/web-developer. It contains a link to the public W3C validator, allowing you to check the page that your browser currently is viewing.

Some examples of valid HTML (and XHTML) document declarations are at htmlhelp.com/tools/validator/doctype.html.

The W3C validator is at validator.w3.org, and the source code for the validator is at validator.w3.org/source.

The home page for the Tidy library is at tidy.sf.net. The home page for the Tidy gem for Ruby is rubyforge.org/projects/tidy. The home page for the Rails Tidy plugin is at www.cosinux.org/~dam/projects/rails-tidy/doc.

The `html-test` plugin for Ruby on Rails is at github.com/Empact/html_test/tree/master. This project at Github has some documentation, as well as the code itself.

Finally, the book *Advanced Rails Recipes*, edited by Mike Clark and published by the Pragmatic Programmers, has a short recipe (#57) by Peter Marklund describing the use of HTML validation from within Rails automated testing. I generally have found this book to be an excellent source of inspiration, along with informing me of many plugins and gems that I had not yet discovered.

EtherDrive®

The AFFORDABLE Network Storage



Fibre Channel speeds at Ethernet prices!

Is your budget shrinking while your network storage needs are growing? Are you suffering from “sticker shock” induced by expensive Fibre Channel and iSCSI storage area network solutions? EtherDrive® SAN solutions offer Fibre Channel speeds at Ethernet prices! Starting at just \$1,995 for a 4TB system, EtherDrive® is the **affordable** storage area network solution. With sustained access speeds from 200MBytes/sec to over 600MBytes/sec, EtherDrive® SAN solutions are **fast**. From a 4TB single storage appliance to multi-PetaByte system by simply adding more storage appliances, EtherDrive® SAN solutions are **scalable**. From a single storage appliance to a network of sophisticated virtualized storage LUNs, EtherDrive® SAN solutions embrace **virtualization**.

Coupling Ethernet technology with SATA hard disk drives, EtherDrive® SAN solutions exploit commodity components to deliver **affordable, fast** storage area network solutions that keep more green in your wallet! Whether you use your own SATA compliant disk drives or our certified enterprise class disk drives, you are in control! EtherDrive® SAN solutions accept standard SATA hard disk drives. Ethernet and SATA disk drives - two proven technologies in one **affordable, fast** storage area network solution - EtherDrive®.

EtherDrive® SAN solutions use the open ATA-over-Ethernet (AoE) lightweight network storage protocol. Simple. Easy to understand. Easy to use. AoE uses Ethernet to transport ATA disk commands without the burden of TCP/IP overhead, thereby enabling disk drives to become AoE devices connected directly to an Ethernet network. An AoE device can be a single physical disk or a logical device made up of multiple disks. An EtherDrive® SAN appliance is an AoE target device.

Finally, an **affordable, fast** storage area network solution for your VMware® ESX 3.5 installation. The EtherDrive® VMware ESX Host Bus Adapter empowers ESX with AoE technology to deliver EtherDrive® SAN solutions for your VMware ESX 3.5 installation.

Shipping EtherDrive® RAID solutions since 2004, Coraid boasts thousands of satisfied customers spanning a broad spectrum of the market including enterprise, government, educational institutions, and hosting service providers. **Call today** to order your EtherDrive® solution, and join the ranks of our thousands of satisfied customers!

Call 1.877.548.7200
or visit our website at
www.coraid.com
International: +1.706.548.7200



 **vmware®** | technology alliance
PARTNER
ESX 3.5 compatible EtherDrive® HBA



MARCEL GAGNÉ

Serious Cool, Sysadmin Style!

Open-source software users have practically thousands of cool packages from which to choose. It's too many packages and too little time. Desktop users know this joy well. Now system administrators can know the same joy, and then some.

So, mon ami, I've granted your wish, and you've had lots of time to show me the packages you've chosen for today's menu. Let's see what you've come up with. But François, there are literally hundreds of choices here! It would take forever to load, configure and test all these packages. I must admit, you've made some excellent choices, but we can't possibly cover all these things. Don't fret. There may yet be a way, but we'll need the help of another great package to do this right. Good guess, *mon ami*, that is exactly what I am talking about. *Quoi?* Oh, I see. You have that one on your list as well. Perhaps you do, but I may show you a side of things you hadn't thought about.

Quickly, François! Put that list aside and get ready. Our guests will be here momentarily. In fact, I see them coming to the door now.

Welcome, everyone, to *Chez Marcel*, not only one of the world's great restaurants, but also a special dining experience where great open-source software meets great wine—and, of course, great customers. Please, *mes amis*, take your tables, sit down and be comfortable. Tonight's wine selection will arrive shortly, courtesy of my faithful waiter. François, kindly head down to the wine cellar and bring back the 2006 Quinta Do Infantado from Portugal. Henri brought in three cases today and left them by that collection of alien artifacts. *Merci*, François. Given today's topic, *mes amis*, we will be serving something a little different, a deep cherry-chocolate, ruby-red port.

This time around, I've decided to let my faithful waiter select the menu. Unfortunately, as my mother used to say, his eyes are bigger than his stomach. Consequently, we needed to enlist the help of a really cool program to run his selection of really cool programs. As François will attest, putting together a menu can be difficult when so many great packages exist—the small, relatively trouble-free desktop applications you can install and try from your package manager, whether it be Synaptic, YaST or your distribution's favored package manager. To check some of them out, fire up your package manager

and take a look at what's available. It's like Christmas or your birthday every time you look—free software, and lots of it, at your fingertips.

The computer I am using right now is running OpenSUSE, and frankly, it came with lots of cool software. But, as with all good things, I always want more, as would most people in this restaurant, I am sure. To try some cool new package, without knowing what it might be, start YaST and select Software Management from the menu. A second window opens from which you can search for a particular package. Let's say you want to install an instant-messaging client—other than the one your distribution came with, that is. Enter the word “instant” in the search field, and all of the packages that have instant in either their package name or description appear in the window to the right. You suddenly learn about Empathy or Pidgin and decide to give them a try.

Click on a package name, and a description of the software appears in the tabbed Description window in the right lower half of the screen (Figure 1). If this is the package you want, click on the check box next to the package name, then click the accept button in the bottom right-hand corner. Should there be dependencies associated with the package you chose to install, a pop-up window appears informing you of that fact. Click Continue, and the installation proceeds. That's all there is to it. If you'd rather browse, click the drop-down box labeled Filter, and select Package Groups to discover packages arranged according to, yes, groups. Looking through games, say you discover a cool-sounding program called *Ri-Li*, a wooden train arcade game, and decide to install it (Figure 1). By the way, *Ri-Li* actually is a great game, and I highly recommend it. Your kids will love it.

Installing software using the package manager from other distributions is just as easy, even if it does look a little different (see Synaptic's interface in Figure 2). This is all great fun and a great way to discover some of the amazingly cool software that free and open-source programmers have created.

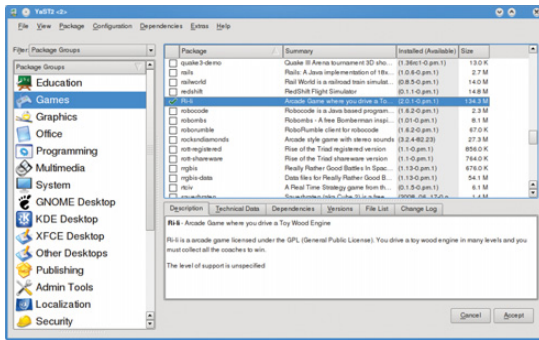


Figure 1. Installing software in a modern distribution like OpenSUSE, is drop-dead simple.

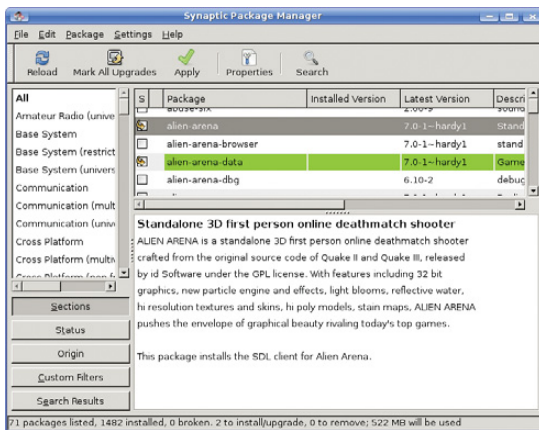


Figure 2. With installers like Synaptic, there's almost no reason not to try out at least one cool package every day.

There's also little to worry about if you do your hunting and installing through your distribution's package manager. With multiple ways to search, descriptions of the packages and automatic installation of prerequisite packages, there's almost no reason not to load up, experiment and discover cool new stuff.

However, if you are looking to try out packages that are more server-oriented, you might be a bit more reluctant. For situations where that cool software effectively comes down to installing a server and all its associated packages, things can become a bit more complex. For instance, installing a content management system isn't just a matter of downloading a package and having the prerequisites install automatically. You may not find the package in your distribution's repository at all. Once you know what you want to try, you'll still need a computer configured with an Apache server, PHP, MySQL (or PostgreSQL), a mail agent like Postfix, a handful of Perl modules and possibly a great deal more. That's why servers often are still the realm of career system administrators (and also why they get the big bucks).

Wouldn't it be awesome if you could just download a server running content management systems like Drupal or Joomla!, or customer relationship management software like SugarCRM or vTiger?

Wouldn't it be awesome if you could just download a server running content management systems like Drupal or Joomla!, or customer relationship management software like SugarCRM or vTiger? Maybe what you really want to do is take an enterprise document management package like Alfresco for a spin or set up a bug-tracking system like Mantis or Bugzilla. And, what if you could have all those rather more complex prerequisites like the Web server, the mail agent and so on already taken care of? Well, you can. Several companies offer prebuilt servers running great open-source software packages like those I've mentioned. You just need to know where to look and how to run them.

Many of these systems are built as VMware images, though not exclusively. You'll find images to run on QEMU or KVM (both of which I've covered in earlier Cooking with Linux columns), Parallels, VMware and others. All of these packages perform hardware virtualization, literally reproducing a PC's hardware in memory so that you can install and run other Linux distributions (or BSD or Windows) on your PC. You create a virtual disk, boot from a CD or CD image, and install on that virtual hard drive. Then, you can run that new machine on your current desktop.

I am currently running OpenSUSE on this notebook, but I have several virtual machines installed as well. In a few seconds, I can start a virtual machine running Mandriva, Fedora, Puppy Linux, Kubuntu, CentOS and others. I do this regularly to test and run different Linux distributions. Those distributions each reside inside a disk image on my system. When the virtual machine is shut down, those distributions and their machines are just big files. That's the idea behind running enterprise software in a virtual appliance and where we start exploring.

Let's start our adventure at the aptly named Virtual Appliances (virtualappliances.net), a company that produces small Linux-based appliances that can be run from a virtual machine. These machines are prebuilt and configured with tools like Cacti, ntop or a LAMP (or LAPP) environment,

and more. Just download, extract, and get ready for some quick machine deployments. For my example, I've decided to download the ntop appliance in VMware format. Because I don't have VMware on my notebook, I took advantage of VMware's free VMware Player, available from www.vmware.com/download/player. This is not the full VMware virtualization suite from which you can install and build your own machine. It is literally a player, as though that virtual appliance you are downloading were a movie you wanted to watch—not just any movie, but a really cool movie you can interact with.

First, download your machine from Virtual Appliances, and extract the tarred or zipped bundle somewhere on your hard disk. Next, download and install VMware Player from the site—you'll find versions for a number of architectures. When you start VMware Player (Figure 3), it offers some basic options that get right to the heart of the matter. You can open an existing virtual machine or download a virtual appliance. Click the Open button and navigate to where you have extracted the virtual appliance, then boot it.



Figure 3. The VMware Player lets you play a virtual appliance like you would play a movie—a movie you can interact with no less.

Once the machine is booted, a message tells you the address you can use to log on to the VA Management Console, in this case <https://192.168.1.109:8000> (Figure 4). Make sure you read the final boot messages so you can get the right address. Open your favorite browser, surf to this address, and enter the console's user name and password (admin and admin). From here, you'll be able to configure the virtual machine further or get information on the various packages that are installed. For instance, the VA console tells me that ntop is running on http port 3000.

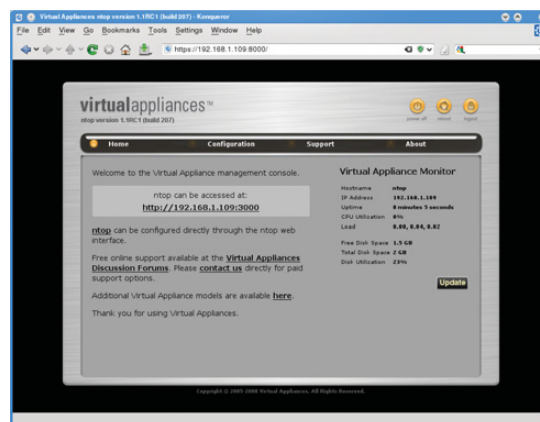


Figure 4. Logging In to the VA Management Console

Using the information provided, I then can start using the installed software. ntop now sits on my system, listening to network traffic and gathering statistics (Figure 5). Everything about this feels like I am running a separate machine. It has its own IP address, runs independently of any other system on the network and is self-contained.

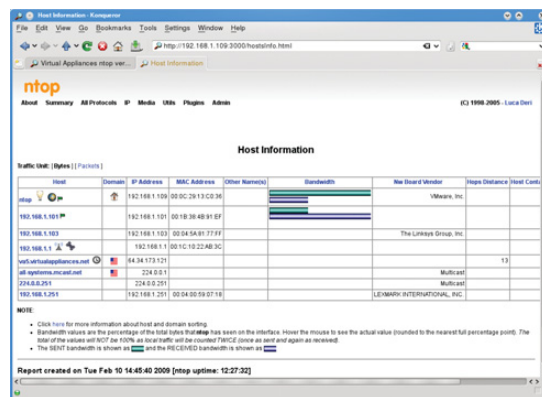


Figure 5. The virtual appliance, in this case using ntop, is up and rolling seconds after the boot completes.

Before I move on, remember that Download button on the front of the VMware Player? That button will open a browser to VMware's collection of virtual appliances, many of which are free, community-contributed builds. It's also a great place to look for other virtual appliances. There's a huge selection sorted into categories along with descriptions and user ratings.

Another site you might want to visit is JumpBox (www.jumpbox.com). Once again, there are tons of virtual machines available, sorted into categories. JumpBox builds machines running the latest enterprise applications, but it does charge for this service (although at \$149 annually, it seems inexpensive). JumpBox does, however,

provide slightly older releases for free. Even if you don't want to shell out the dollars for a membership, you still can download and evaluate a number of great packages.

The VMware Player isn't the only game in town. Another great piece of virtualization software is VirtualBox, an open-source package freely distributed under the GPL. It's one I use every day, and one I highly recommend. Let's use VirtualBox to run an appliance from JumpBox. I've selected and downloaded a free copy of SugarCRM for this demonstration.

Visit virtualbox.org, download a copy of VirtualBox for your distribution, and install it. When you start VirtualBox the first time, there are no machines running in it. Think of it as a blank slate, or better yet, a new computer with a blank hard drive waiting for your favorite distribution (Figure 6).

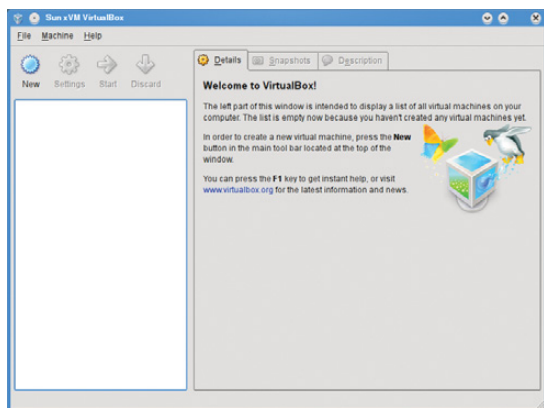


Figure 6. First time out with VirtualBox—think of it as a white-box PC with no OS.

Next, you need to tell VirtualBox about the virtual appliance image. To do this, click File on the menu bar and select Virtual Media Manager. When the window appears (Figure 7), you can start adding the virtual disk images from which you'll boot your machine. Click the Add button, then navigate to the SugarCRM virtual appliance folder. Look for the root folder and attach the root.vmdk file. Usually, that vmdk file is all you need, but with JumpBox, there's another step that I'll visit shortly. Click OK to continue.

Unlike VMware Player, VirtualBox is the whole application, which means you can create different configurations of virtual machines, make a virtual hard drive and install a brand-new machine onto that disk. Click the New button on the top right, and you are presented with a wizard that takes you through all the steps necessary to create this machine. The first step is to name this machine and tell VirtualBox what OS it will be running (Figure 8). Click Next, and VirtualBox asks you




Figure 7. You'll need to tell VirtualBox about your virtual appliance so the machine you create can boot it.

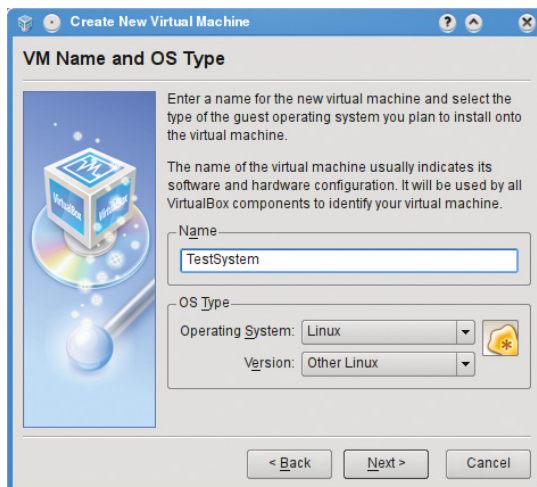


Figure 8. The wizard guides you through the steps for creating a virtual machine.

how much memory (RAM) you want to give this machine. The default is 256MB. Click Next again, and you're asked about the hard disk you want to use.

This is where things get interesting. If you choose to use an existing disk, from an existing virtual machine, you can select it from the drop-down list. Machines you added from the Virtual Media Manager will appear here. On another day, you would click New and create a hard drive onto which to load the latest Ubuntu, Mandriva or whatever your favorite distribution might be. Assuming you went the virtual appliance route, selecting the image name, then click Next and you're almost done. Your new virtual machine is listed in the left sidebar (Figure 9).

One last thing, and this is where it's actually easier with VMware's Player (which isn't GPL'd software, unfortunately). You'll see only one

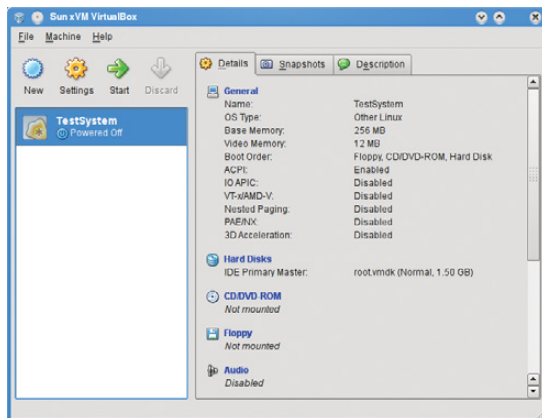


Figure 9. Once a new machine has been created in VirtualBox, it shows up in a list on the left sidebar. There can be many machines here.

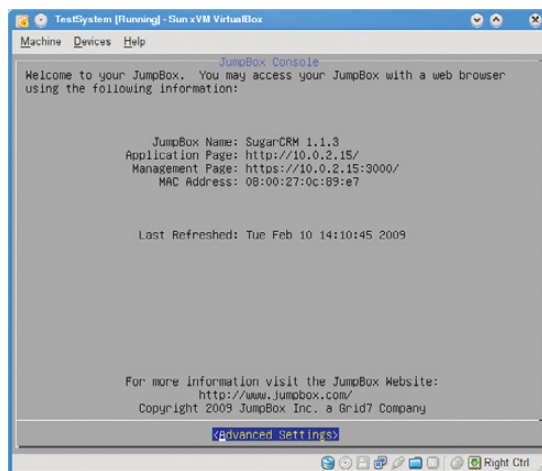


Figure 10. Your SugarCRM JumpBox is up and running under VirtualBox. Follow the links provided on the boot page to access your new machine's application or management page.

hard disk attached on the left. JumpBox appliances generally use two virtual disks for each machine: one for the root (`root.vmdk`) and one for data (`data.vmdk`). You need to add the data disk as well. Click the blue Hard Disks link, then navigate to the data disk and add it. The only thing you really need to be careful about here is making sure the root disk is first in line, as VirtualBox will boot from the hard disk. You'll find yourself back at the VirtualBox start screen but with at least one virtual machine ready to start. Click the Start icon, and your virtual appliance boots. Once booted, the virtual machine displays some information about the machine. On first boot into a virtual appliance, that screen most likely will have three links (Figure 10).

One link will take you to a page where you can

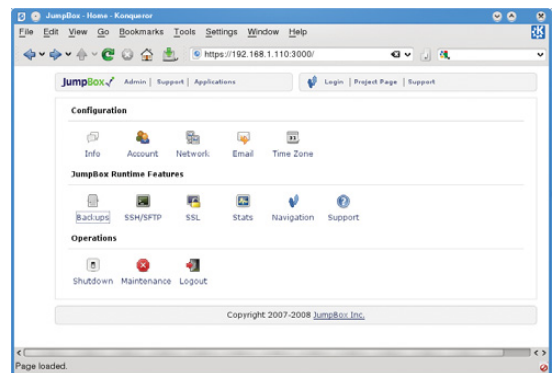


Figure 11. JumpBox virtual appliances come with a great Web-based administration console.

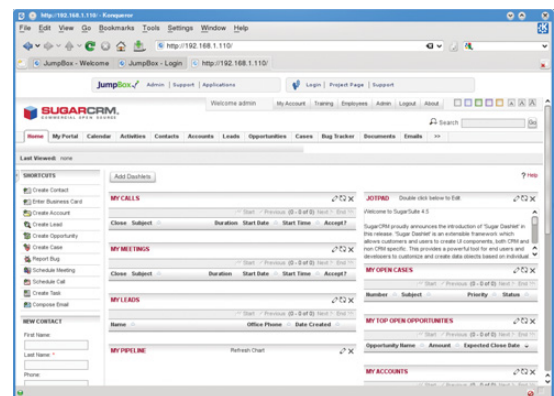


Figure 12. What it's all about: a virtual appliance running enterprise-class software like SugarCRM in no time at all.

finish configuring your machine—usually a minor task as almost everything else is done for you in the virtual appliance. There also will be links to access the application's page and its administration console. The JumpBox administration page gives you access to basic machine operations, such as performing a shutdown or running a backup so you can recover the machine state should disaster strike (Figure 11).

Of course, the real excitement comes from trying out that cool application or suite. By downloading a JumpBox virtual appliance and simply booting it (in either VMware Player or VirtualBox), I pretty much have instant access to a full SugarCRM implementation without all those steps involving Web servers, databases and so on (Figure 12).

The new virtual machine runs like any other machine, and in some ways, it runs better. You can turn off a virtual machine and save its execution state so that when you reboot, at a later time, everything is exactly as it was. Any open application

is open as it was. This kind of technology—the ability to load up virtual appliances and deploy them in minutes—is what cool really means. Take some time to check out Virtual Appliances, JumpBox and VMware’s Virtual Appliance Marketplace, and I guarantee that it will change your sysadmin life forever.

As you can see, *mes amis*, it is possible to have it all, at least in a virtual sense. Best of all, you can have it fast (after the download completes). Unfortunately, we cannot save the restaurant’s current state or that of the wine. All open bottles must be emptied; a delightful imposition, I am sure you will all agree. François, please attend to our guests and refill those glasses once more before we say *Au revoir*. Please, *mes amis*, raise your glasses and let us all drink to one another’s health. *A votre santé! Bon appétit!* ■

Marcel Gagné is an award-winning writer living in Waterloo, Ontario. He is the author of the *Moving to Linux* series of books from Addison-Wesley. 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 marcel@marcelgagne.com. You can discover lots of other things (including great Wine links) from his Web sites at www.marcelgagne.com and www.cookingwithlinux.com.

Resources

JumpBox: jumpbox.com

Virtual Appliances: virtualappliances.net

VirtualBox XM: www.virtualbox.org

Download Free VMware Player:
www.vmware.com/download/player/download.html

VMware’s Virtual Appliance Marketplace:
www.vmware.com/appliances

Marcel’s Web Site: www.marcelgagne.com

Cooking with Linux: www.cookingwithlinux.com

WFTL Bytes!: wftlbytes.com

WHICH ONE ARE YOU?

THE BAD GUYS KNOW YOUR TRICKS. DO YOU KNOW THEIRS?

Are you the kind of security professional that craves the freshest information?
 Are you the kind of hands-on network defender that needs to be the first to know?
 If you stand out from the crowd because you’re seriously committed to security,
 then you’re our type – the Black Hat type.

The Black Hat Briefings are the most important technical security events on the planet, and we create them for people just like you – people who know security and want to be part of the solution. Join us in Las Vegas for 10 full tracks and over 60 intense hands-on training sessions of timely, actionable security information.

Black Hat

BRIEFINGS AND TRAINING

SPONSORS

DIAMOND

Microsoft

PLATINUM

CISCO

RSA
The Security Division of EMC

QUALYS

GOLD

COMPE IBM

IOActive infosecurity

NORMAN

SILVER

Aladdin
SECURE THE BLACK HAT

CENZIC

NETWISSE

splunk

ArSight

hp

PROLEXIC

BStillSecure

BIG FIX

intel

redseal

Sunbelt Software

Blue Coat

Learning Glass Systems

SAINT

TRUSTED

Booz | Allen | Hamilton

McAfee

SecureWorks

TipplingPoint

BREACH

VE

SOLERA

WHITEHAT



DAVE TAYLOR

Deal or No Deal!

Use bash to calculate game outcomes and help you win at *Deal or No Deal*.

I admit it, my kids and I often find ourselves at the local arcade playing video games. There seems to be two basic types of arcades nowadays: those where you play to earn free games and get the high score, and those that are ticket-based where doing well on a game produces tickets that you cash in for cheesy tchotchkes or incredibly cheap stuffed animals.

We tend to go to the latter, and one of the games that's caught my attention is a video version of the popular TV show *Deal or No Deal*. It got me thinking about odds and probability and how these games actually work.

Deal or No Deal, the Game

The game is incredibly simple: you're told that there is a range of prizes hidden in briefcases, and you choose which to eliminate and win what's left in the last briefcase. It's not too exciting, except as you go along, you're offered the opportunity to take a specific cash prize instead of continuing the game.

On the TV show, there are 26 cases, and the prizes range from \$0.01 to \$1,000,000.00 in uneven jumps. Attractive models hold each case, and the game typically proceeds where the player

game? Even though it's a bit lower than the average expected payout, I'd take it. Deal!

Let's say we're most of the way through the game though, and 20 of the possible prizes are knocked out. What's left are 0.01, 50, 300, 1,000 and 250,000. Your chance of picking the 250,000 prize case? One in five.

The expected payout is \$50,270, and if that's what you're offered, it's significantly better than four of the five possible outcomes you face. My advice? Deal!

It turns out that there's some sort of random factor that's thrown in to the "deal", so in the game itself, they vary up or down a percentile value.

Algorithms and Coding

To experiment with this as a shell script (yeah, it only took me half the column to get to my first line of code), we need to work with arrays—something that's frankly a bit confusing. Here's how we could define the 26 briefcases:

```
declare -a cases=(0.01 1 5 10 25 50 75 100 200 300 400 500 750
1000 5000 10000 25000 50000 75000 100000 200000 300000 400000
500000 750000 1000000)
```

This might be the first time you've seen the `declare` statement in the shell. In this case, the `-a` flag declares the variable as being an array. Its usage, however, is not mandatory—arrays can be implicitly declared simply by using the array assignment syntax (a list of values inside parentheses).

You reference individual array elements with `$var[index]`, but there's a twist because of how the shell parses content. What you need to do is actually wrap it with curly braces: `${var[index]}`. Add a `#`, and you get the number of elements in the array, like this:

```
echo Total number of array elements is ${#cases[*]}
```

The value that'll be printed is 26, just what we want for *Deal or No Deal*. To see value #11, you could use `${cases[11]}`, but that's wrong. Why? Because shell arrays are indexed starting at zero, so case #11 is actually `${cases[10]}`, which, yeah, is pretty confusing.

Let's start by writing the snippet that can calculate expected payout before you have picked a single briefcase out of the collection:

```
for (( val=0 ; $val < ${#cases[*]} ; val+=1 )) ; do
```

This might be the first time you've seen the `declare` statement in the shell.

guesses a bunch of cases to eliminate, in clusters. If the eliminated cases have low-value prizes, that's good—there is, therefore, a better chance of winning big bucks. If they have high-value prizes inside, well, that's not good.

After each set of cases is eliminated, the house offers the player a "deal" to stop playing the game—a payoff that's worth more than the lowest unidentified prize but obviously lower than the highest unidentified prize.

The prizes are 0.01, 1, 5, 10, 25, 50, 75, 100, 200, 300, 400, 500 and 750, and then the big values: 1,000, 5,000, 10,000, 25,000, 50,000, 75,000, 100,000, 200,000, 300,000, 400,000, 500,000, 750,000 and 1,000,000.

Calculating the average payout, assuming you have an equal chance of picking each of the cases at the beginning of the game, is easy: add them up and divide by 26. The result is 131,477.54. If they offered me \$130k instead of playing the

```

sum=$(( $sum+${cases[$val]} ))
done
echo sum = $sum, payout ==$(( $sum / ${#cases[*]} ))

```

The bad news? It turns out that the 0.01 value really messes things up, as we can't do integer math with a non-integer value. So, because the 0.01 value actually proves not to influence things much, I'm going to replace it with the value 0. Here's the output:

```
sum = 3418416, payout = 131477
```

That's reasonably accurate. Our earlier calculation was 131,477.54. Close enough for gaming work!

Now, let's randomly pick out 22 of the cases, calculate the expected payout and offer a "deal" versus the values still remaining.

First, pick out some cases:

```

for (( picked=1 ; $picked <= 22 ; )) ; do
    pick=$(( RANDOM % 26 ))
    if [ ${cases[$pick]} -ne -1 ] ; then
        cases[$pick]=-1
        picked=$(( $picked + 1 ))
    fi
done

```

That gets all but four cases out of the game by setting their value to -1 (remembering that we're using 0 to represent \$0.01). To see what's left in the game and calculate the expected payout of only the remaining values, do this:

```

for (( val=0 ; $val < ${#cases[*]} ; val+=1 )) ; do
    if [ ${cases[$val]} -ne -1 ] ; then
        echo \((Still in the game: a prize \
            worth \${cases[$val]})\
        sum=$(( $sum+${cases[$val]} ))
        cnt=$(( $cnt+1 ))
    fi
done

```

```

echo Win \${(( $sum / $cnt ))} if you stop. \
    Deal, or no deal?
exit 0

```

Run the code, and here's a sample output:

```

(Still in the game: a prize worth $50)
(Still in the game: a prize worth $1000)
(Still in the game: a prize worth $5000)
(Still in the game: a prize worth $400000)
Win $101512 if you stop. Deal, or no deal?

```

Would you take the deal, or try for the \$400,000 case? ■

Dave Taylor has been involved with UNIX since he first logged in to the ARPAnet in 1980. That means that, yes, he's coming up to the 30-year mark now. You can find him just about everywhere on-line, but start here: www.DaveTaylorOnline.com.

Low Cost Panel PC

PPC-E7

- Cirrus ARM9 200MHz CPU
- 3 Serial Ports & SPI
- Open Frame Design
- 3 USB 2.0 Host Ports
- 10/100 BaseT Ethernet
- SSC-12S Audio Interface
- SD/MMC Flash Card Interface
- Battery Backed Real Time Clock
- Up to 64 MB Flash & 128 MB RAM
- Linux with Eclipse IDE or WinCE 6.0
- JTAG for Debugging with Real-Time Trace
- WVGA (800 x 480) Resolution with 2D Accelerated Video
- Four 12-Bit A/Ds, Two 16-Bit & One 32-Bit Timer/Counters





2.6 Kernel

Setting up a Panel PC can be a puzzling experience. However, the PPC-E7 Compact Panel PC comes ready to run with the Operating System installed on Flash Disk. Apply power and watch either the Linux X Windows or the Windows CE User Interface appear on the vivid color LCD. Interact with the PPC-E7 using the responsive integrated touch-screen. 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 \$495.

For more info visit: www.emacinc.com/panel_pc/ppc_e7.htm

Since 1985
OVER
23
YEARS OF
SINGLE BOARD
SOLUTIONS

EMAC, inc.

EQUIPMENT MONITOR AND CONTROL

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



KYLE RANKIN

Lightning Hacks Strike Twice

Are four hacks better than one? This month, check out four short Linux tips that will save you time and effort (or just rotate your X display like a top).

It was exactly one year ago when I wrote my first Lightning Hacks column. The column was inspired by lightning talks that often occur at conferences. In a lightning talk, instead of one speaker giving a presentation for an hour, different people give 5–10-minute presentations. The idea is that many people are working on cool projects but may not have an hour's worth of material to present. And, the audience gets a rapid-fire presentation of a few different topics instead of one long lecture.

The idea behind lightning hacks is similar—I can cover some quick hacks I think are interesting but that don't warrant a full column. For example, in the first Lightning Hacks column [June 2008], I talked about an expanded `wmctrl` script that reset all of my windows to default locations and sizes, another script that toggled my laptop output for when I connect to a projector, and finally, I discussed how to use `rdiff` to create small diff files for large binaries. Now that a year has passed, I think it's time for lightning to strike twice.

Change to Your Previous Directory

Here's a quick one. I've mentioned this trick to a number of people, and I get one of two responses. This is one of those tricks (like `Ctrl-R` shell expansion in bash) that you either already know about and seems obvious to you, or one that you can't believe took so long to discover.

If you have spent a lot of time on the command line, you probably have heard about the `pushd` and `popd` scripts. These scripts let you create a stack that you can push directories on to and later pop

If you have spent a lot of time on the command line, you probably have heard about the `pushd` and `popd` scripts.

them when you want to return to a previous directory. This script is cool, except you have to know in advance you want to save a directory and push it, so you can pop it later—I never seem to remember. Generally speaking, what I need is some quick way

to go back to my previous directory. Lucky for me, bash's `cd` has this feature built in. All I do is type:

```
$ cd -
```

Bash keeps track of your current working directory in the `$PWD` variable and your previous directory in `$OLDPWD`. If you type `cd -`, bash substitutes `-` with `$OLDPWD`. Although you certainly could just type `cd $OLDPWD`, `cd -` is faster and easier to remember.

SSH Key One-Liner

If you have to manage a lot of servers or run remote scripts in cron, SSH keys are a lifesaver. It's so nice to be able to `ssh` to a machine and instantly log in without typing a password. Of course, one of the more annoying parts of the process can be setting up the SSH keys on the remote host. Typically, the process goes something like this: run `ssh-keygen` locally, `scp` the `~/.ssh/id_rsa.pub` to the remote server, then `ssh` to the remote server and append that key to your remote `~/.ssh/authorized_keys` file.

The above method works, but if you can do the entire thing with a one-liner, why wouldn't you? Here's the SSH one-liner that will copy your local SSH key to the remote host, so you have to type the password only once in the whole process:

```
$ ssh user@server.example.net "cat >> ~/.ssh/authorized_keys"
  << ~/.ssh/id_rsa.pub
```

Image Drives over SSH

Many great imaging tools are available, but for me, it's still hard to beat `dd` (unless your drive is dying, in which case you should use `ddrescue`). It is such a powerful, blunt, ancient UNIX tool, it's hard not to love it. These days, I don't image too many drives. I use `kickstart` for server deployments and `rsync` when I want to migrate files. That said, I still do image drives when I want to perform forensics on the host.

One problem you often have when you image drives is that your server might be in a data center hundreds or thousands of miles away. Even if the server is close by, you might not be able to add an extra

WHERE
will YOU be on June 2?



JavaOneSM

JavaOneSM Conference | June 2–5, 2009

The Pavilion: June 1–4, 2009, The Moscone Center, San Francisco, CA

The JavaOneSM conference brings together developers, technology enthusiasts, and industry luminaries from around the world. It's your chance to learn, grow, and network with the vast—and growing—open technology community.

This year's JavaOne conference offers even more opportunity to grow your skills. You can:

- Connect with developers from more than 70 countries
- Hear from expert speakers on the hot topics you care about most
- Choose from a wide variety of targeted tracks, labs, and BOFs
- Get tips and best practices from technology creators and evangelists for areas such as next-generation Web services and cloud platforms
- Experience groundbreaking technologies—hands-on in our Pavilion

Nowadays money's tight. That's why it's more important than ever to attend the one conference that delivers everything you want to see, learn, and experience — all under one big roof. And that's the JavaOne conference.

Save Your Spot—Register Today!

By the way, check out CommunityOne, Sun's conference on open-source innovation and implementation, colocated with the JavaOne conference. developers.sun.com/events/communityone

Save \$100

on Conference registration!

Register by June 1st at
java.sun.com/javaone

Platinum Cosponsor



Cosponsors



Sony Ericsson



© 2009. All rights reserved. Sun, Sun Microsystems, the Sun logo, Java, JavaFX, and JavaOne are trademarks or registered trademarks of Sun Microsystems, Inc. or its subsidiaries in the United States and other countries. Information subject to change without notice.

drive on the fly. In either case, most sysadmins end up imaging the drive over the network. Traditionally, this was done via netcat, but these days, you always have to figure out some port you can use that won't be blocked by the firewall. Another problem is that netcat will transmit potentially sensitive data over the network in plain text. The modern solution to this problem is to use SSH. Many servers now have SSH running and available out of the box, and with modern processor speeds, the encryption overhead shouldn't be too bad either.

The one-liner to image drives over SSH works much like the one I used for SSH keys above. It takes advantage of the fact that if you pipe data or redirect input to SSH on the command line, it will forward it to the remote connection. So, if I wanted

to image /dev/sda on my local machine to a file called /media/disk1/sda-image.img on server.example.net, I would type the following:

```
$ sudo dd if=/dev/sda | ssh user@server.example.net "cat
➔ /media/disk1/sda-image.img"
```

If I didn't want to image to a file and instead wanted to image directly to a drive on server.example.net, I simply could replace /media/disk1/sda-image.img with that device file (I just would need to log in as root).

Because you can image a drive over SSH, it makes sense that you can use a variation on the command to restore your image back to a drive. Here's the inverted version of the above command that I would use if I wanted to restore the /media/disk1/sda-image.img image I created back to /dev/sda:

```
$ ssh user@server.example.net "cat /media/disk1/sda-image.img"
➔ | sudo dd of=/dev/sda
```

Rotate Your Screen Around and Around

My laptop doubles as a tablet, and even though I don't use the tablet mode very often, when I do use it, I like to be able to rotate the screen around to portrait mode and back. Now, dock applications exist that can do this with a few clicks, and I always could just try to remember the right xrandr commands, but instead, I wrote a little script that I then bound to one of the hardware buttons on my laptop display. Each time I press the button, it runs the script and rotates the screen another 90 degrees.

The key to the script is to keep track of your current orientation. When xrandr rotates, it rotates only left, right, inverted or normal, so if you already are rotated to the left and rotate left again, it won't change. To accomplish this, I just write the current orientation to a temporary file. Listing 1 shows the full script.

Notice in Listing 1 that I also added an echo piped to osd_cat. This is optional and just displays the current orientation on my screen. If you want to use this, be sure you have the osd_cat utility (it's included with the xosd-bin package in Debian and Ubuntu). The way the script is set up, it will run through each of the orientations in order before it goes back to normal. Because the temporary file will be deleted any time the machine reboots, I made sure to set the default mode to rotate 90 degrees. ■

Kyle Rankin is a Senior Systems Administrator in the San Francisco Bay Area and the author of a number of books, including *Knoppix Hacks* and *Ubuntu Hacks* for O'Reilly Media. He is currently the president of the North Bay Linux Users' Group.

Listing 1. Screen Rotation Script

```
#!/bin/sh

export ORIENTATION=`cat /tmp/.orientation`

if [ $ORIENTATION -eq "90" ]; then
    xrandr --auto
    xrandr --output LVDS --rotate inverted
    echo 180 > /tmp/.orientation
    echo "180" | osd_cat --shadow=2 --align=center \
    --pos=bottom --color=green --delay=2 \
    --font=lucidasanstypewriter-bold-24 \
    --offset 40 &
elif [ $ORIENTATION -eq "180" ]; then
    xrandr --auto
    xrandr --output LVDS --rotate left
    echo 270 > /tmp/.orientation
    echo "270" | osd_cat --shadow=2 --align=center \
    --pos=bottom --color=green --delay=2 \
    --font=lucidasanstypewriter-bold-24 \
    --offset 40 &
elif [ $ORIENTATION -eq "270" ]; then
    xrandr --output LVDS --rotate normal
    echo "Normal" | osd_cat --shadow=2 --align=center \
    --pos=bottom --color=green --delay=2 \
    --font=lucidasanstypewriter-bold-24 \
    --offset 40 &
    echo 0 > /tmp/.orientation
else
    xrandr --auto
    xrandr --output LVDS --rotate right
    echo 90 > /tmp/.orientation
    echo "90" | osd_cat --shadow=2 --align=center \
    --pos=bottom --color=green --delay=2 \
    --font=lucidasanstypewriter-bold-24 \
    --offset 40 &
fi
```

VMWARE CERTIFIED SOLUTIONS

ABERDEEN STIRLING 123



1U Dual Xeon VMware Certified Server

- Up to two Dual-Core or Quad-Core Intel® Xeon® processors
- Up to 32GB 667/533MHz Fully Buffered ECC DDR2 SDRAM
- Up to 4 x Hot-Swap SATA or SAS Hard Drives
- Universal I/O allows for 3 expansion cards in 1U
- Pre-installed VMware® ESXi on Disk-on-Module
- 650W High-efficiency Redundant Power Supply
- 5-Year Warranty

Starting at **\$1,999**

ABERDEEN STIRLING 229



2U Dual Xeon VMware Certified Server

- Up to two Dual-Core or Quad-Core Intel Xeon processors
- Up to 128GB 800/667/533MHz Fully Buffered ECC DDR2 SDRAM
- Up to 8 x 1TB (8.0TB) Hot-Swap SATA Hard Drives
- Up to 7 x Low-Profile Expansion Slots
- Pre-installed VMware ESXi on Disk-on-Module
- 700W High-efficiency Redundant Power Supply
- 5-Year Warranty

Starting at **\$2,425**

ABERDEEN STIRLING 428T



4U Dual Xeon VMware Certified Server

- Up to two Dual-Core or Quad-Core Intel Xeon processors
- Up to 64GB 667/533MHz Fully Buffered ECC DDR2 SDRAM
- Up to 8 x 1TB (8.0TB) Hot-Swap SATA Hard Drives
- Up to 6 x Full Height Expansion Slots
- Pre-installed VMware ESXi on Disk-on-Module
- 800W High-efficiency Redundant Power Supply
- 5-Year Warranty

Starting at **\$2,350**

ABERDEEN STIRLING 132T



1U Twin Node VMware Certified Server

- Up to two Dual-Core or Quad-Core Intel Xeon processors/node
- Twin Nodes allows for up to 4 processors & 16 cores in 1U
- Up to 64GB 800/667/533MHz Fully Buffered ECC DDR2/node
- Up to 2 x 1TB Hot-Swap SATA Hard Drives per node
- Pre-installed VMware ESXi on Disk-on-Module
- 980W High-efficiency Power Supply
- 5-Year Warranty

Starting at **\$2,880**

ABERDEEN STIRLING 244



2U Quad Xeon MP VMware Certified Server

- Up to four Quad-Core or Six-Core Intel Xeon MP processors
- Quad Six-Core allows for 24 processor cores in 2U
- Up to 192GB 667/533MHz Fully Buffered ECC DDR2 SDRAM
- Up to 6 x Hot-Swap SATA or SAS Hard Drives
- Pre-installed VMware ESXi on Disk-on-Module
- 1200W High-efficiency Redundant Power Supply
- 5-Year Warranty

Starting at **\$5,615**

ABERDEEN STIRLING 444



4U Quad Xeon MP VMware Certified Server

- Up to four Quad-Core or Six-Core Intel Xeon MP processors
- Quad Six-Core allows for 24 processor cores in 4U
- Up to 192GB 667/533MHz Fully Buffered ECC DDR2 SDRAM
- Up to 5 x Hot-Swap SATA or SAS Hard Drives
- Pre-installed VMware ESXi on Disk-on-Module
- 1200W High-efficiency Redundant Power Supply
- 5-Year Warranty

Starting at **\$6,625**

ABERDEEN ISCSI XDAS



DAS VMware Certified Expandable Storage

- IP SAN Solution
- Single or Redundant Controller
- Expandable up to 64TB in a single array
- 2U/12 Bay and 3U/16 Bay Models available
- SAS or SATA Hard Drive Support
- Fault-tolerant Modular Hardware Design
- 5-Year Warranty

Starting at **\$7,995**

ABERDEEN FIBRE XDAS



DAS VMware Certified Expandable Storage

- Hardware RAID5 and RAID6 engine by dedicated ASIC400
- Single or Redundant Controller
- Expandable up to 64TB in a single array
- 2U/12 Bay and 3U/16 Bay Models available
- SAS or SATA Hard Drive Support
- Fault-tolerant Modular Hardware Design
- 5-Year Warranty


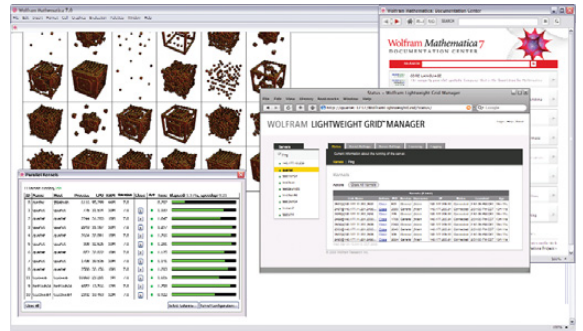
Starting at **\$8,495**



Wolfram Research's gridMathematica

Wolfram Research's new gridMathematica 7 enables users to utilize the built-in parallelization capabilities of its Mathematica application and, thus, run more tasks in parallel on more powerful hardware and clusters. gridMathematica adds extra computation kernels and automated network distribution tools, allowing users to achieve faster execution "without changing a line of code", says Wolfram. Three different products are part of the series: gridMathematica Local, gridMathematica Server and Wolfram Lightweight Grid Manager. gridMathematica requires Mathematica and is available for Linux, Mac OS X and Windows.

www.wolfram.com/products/gridmathematica



John Lamb's *The Greening of IT* (IBM Press)

More so than nearly all its rivals, IBM has made "going green" a core mission. Not only has IBM rolled out its "Big Green" and "Big Green Linux" initiatives, but it also has now published one of the few books on green IT, called *The Greening of IT: How Companies Can Make a Difference for the Environment*. In the book, IBM senior staffer, John Lamb, tackles both macro and micro issues surrounding the reduction of the environmental impact caused by IT operations. At the macro scale, Lamb looks at the role of governments and electrical utilities and the importance of good regulations and incentives. At the micro level, Lamb examines the nuts and bolts of reducing energy consumption in the data center, covering organizational issues, ROI, procurement, asset disposal, measurement of energy consumption, virtualization, cooling equipment and much more. Finally, the author explores case studies of all types and sizes worldwide, including IBM's own \$1 billion Big Green initiative.

www.ibmpressbooks.com

Super Talent's UltraDrive Family of Solid-State Drives

The crew at Super Talent has been busy preparing not one but two new families of solid-state drives (SSDs), the UltraDrive ME and UltraDrive LE. The company calls the lines "next-generation SSDs" that offer "noticeable performance gains at boot time, application loading and accessing data". Although both lines offer 32GB, 64GB and 128GB variants, the UltraDrive ME line offers an additional 256GB model. The UltraDrive LE is rated for a maximum sequential read speed of 230MB/s, while the UltraDrive ME comes in at 200MB/s. Regarding maximum sequential write speed, the UltraDrive LE clocks 170MB/s, and the UltraDrive ME at 160MB/s. Super Talent says that the drives are designed to be "compatible with all known operating systems", including Linux, DOS and Windows.

www.supertalent.com



ScaleMPTM

ScaleMP's vSMP Foundation

Making the area of virtualization even more interesting is ScaleMP's updated Versatile SMP (vSMP) Foundation 2.0 virtualization solution. vSMP Foundation aggregates multiple industry-standard off-the-shelf x86 servers (rackmounted or blade systems) into one single virtual high-end system for the HPC market. This new release of vSMP, says ScaleMP, offers "significantly enhanced performance" through support for the forthcoming Intel Nehalem processor family, as well as enhanced enterprise-class features, such as increased high-availability, partitioning of a single virtual system into multiple isolated environments, extended remote management, enhanced profiling capabilities and support for Emulex LightPulse Fibre Channel HBAs.

www.scalemp.com



Compiere Cloud Edition ERP

Compiere ERP—a comprehensive open-source application that automates business processes, such as accounting, purchasing, order fulfillment, manufacturing, warehousing and CRM—is now available on the Amazon Elastic Compute Cloud (EC2). The new Compiere Cloud Edition is delivered with a complete technology stack—that is, an operating system, application server and database that can be deployed on Amazon EC2 “in a matter of minutes”. Compiere says that the “convenient virtual computing environment” reduces the cost of ERP deployment by eliminating up-front capital costs for hardware and software and reducing ongoing IT infrastructure support costs. The company also notes the advantages of cloud computing, which allows IT departments to increase capacity or add capabilities “on the fly” without investing in new hardware, personnel or software by accessing virtual servers available over the Internet to handle computing needs. A range of subscriptions include application support, service packs and access to Compiere automated upgrade tools.

www.compiere.com

Justin Seitz's *Gray Hat Python* (No Starch Press)

Publisher No Starch Press touts Justin Seitz's new book *Gray Hat Python* as “the first Python book written for security analysts”. Subtitled “Python Programming for Hackers and Reverse Engineers”, the book explains the intricacies of using Python to assist in security analysis tasks, teaching readers how to design debuggers, create powerful fuzzers, utilize open-source libraries to automate tedious tasks, interface with security tools and more. *Gray Hat Python*, says No Starch, covers everything from the nuts and bolts of how to use the language for basic code and DLL injection to using Python to analyze binaries and disassemble software. More than anything, however, the book reveals how superior the Python language is when it comes to hacking, reverse engineering, malware analysis and software testing.

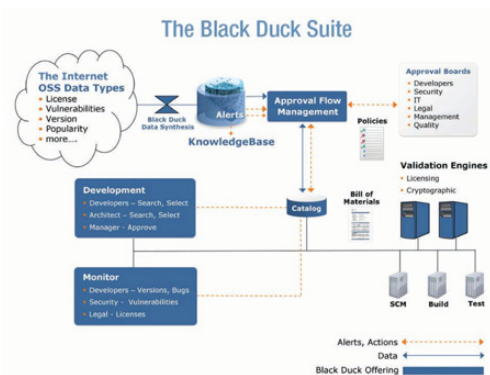
www.nostarch.com



Black Duck Software's Black Duck Suite

The gist behind Black Duck Software's new Black Duck Suite is to give development organizations a comprehensive management platform for taking advantage of open-source components while addressing the associated management, compliance and security challenges. Black Duck says that its new product brings “new levels of automation and efficiency” to these tasks and “enables developers to focus on creating innovative business value instead of ‘re-inventing the wheel’”. Black Duck Suite is a unified framework of the company's Code Center, Export and Protex enterprise products, plus SDK with Web services API that integrates with other tools and environments. Key product features include a searchable internal catalog, a customizable approval work flow and a comprehensive KnowledgeBase of open-source information.

www.blackducksoftware.com



BlueStripe Software's FactFinder

Rounding out the trio of memorable color + object company names is BlueStripe Software, which recently released version 2.0 of FactFinder, an application for staging, deploying and managing business-critical applications. Now available for Red Hat Enterprise Linux, FactFinder enables “unsurpassed intelligence into the performance and behavior” of applications, allowing users to understand their structure and relationship to each other, efficiently manage them, identify performance issues and perform triage to resolve issues. Key new features include automatic discovery and mapping, health and performance measurement and service-level driven triage.

www.bluestripe.com



Please send information about releases of Linux-related products to newproducts@linuxjournal.com or New Products c/o Linux Journal, PO Box 980985, Houston, TX 77098. Submissions are edited for length and content.

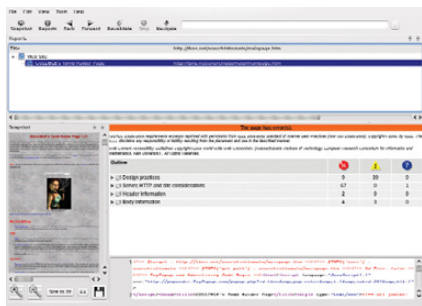
Fresh from the Labs

RightWebPage—Web Conformance Tool

rightwebpage.org

For people designing their own Web pages, editing someone else's or simply learning how to make their own, RightWebPage is an invaluable tool that puts Web pages through the most stringent of tests. Not only does it follow strict W3C standards, it also supports Linux, Mac OS and Windows. According to the Web site:

RightWebPage verifies and corrects many aspects of conformance of Web sites to standards of Internet best practices, such as W3C WCAG and ISO/IEC 23026:2006 - Software Engineering - Recommended Practice for the Internet - Web Site Engineering, Web Site Management and Web Site Life Cycle.



RightWebPage—the Web Conformance Checker that Turns Pedantry into an Art Form

Installation Head to the Web site's download section, and under GNU/Linux, there are the options of RPM and deb packages in both 32- and 64-bit flavors, as well as the obligatory tarball. I went with the deb package and ran into some dependencies. Once I installed the packages `tidy` and `opensp` though, things were fine. I'm ashamed to say that I couldn't work out how to compile the source package though; if you're looking for a `./configure, make, make install`, it's not here. However, that is how it was at the time of this writing, and in the space of a month (the last time I looked), things have changed a great deal, so it may have

improved by the time you read this.

Once you have RightWebPage installed, you can run it with the command:

```
$ rightwebpage
```

Usage As soon as the program starts, you're presented with an introductory splash screen that takes you through a series of prompts having to do with needed information and parameters to check before you get started. First up is the Web page URL (local or Internet). Next is Conformance Tests to run, Document Type, Excessive Scrolling, Accessibility, User Agent Emulation, Character Encoding...at this point, I got bored and clicked Finish. If you choose a Web page to load, for a while no information appears on the screen. Give it a few moments, and eventually everything appears. Oh, and a word of warning, this process can be very CPU- and RAM-intensive.

Once your Web page has loaded, your data is split into four panels. The first covers the filename(s) of what you're looking at. Next are all the faults it has found (and there probably will be a lot of them), which are broken down into categories and can be browsed at will. The bottom panel contains the actual coding of the page you've loaded, which also can be edited and saved and is quite helpfully color-coded. And, the pièce de résistance, on the left is a preview panel of your Web page fully rendered. Any changes you make can be saved, and you can revalidate the Web page at any time.

Now I must warn you, your ego is in for a bruising. RightWebPage is beyond anal in its stringency. I haven't written any Web sites for a number of years, but when I have, I prided myself on strict adherence to Web standards. Well, to my embarrassment, RightWebPage found 80 warnings and 42 alerts in a page that I thought would fare well! So, if you're looking for a strict teacher, this is it. All in all, this is a brilliant piece of software that will keep you on the straight and narrow in Web design, as long as you're willing to submit to

something that is *seriously* pedantic.

Gourmet Recipe Manager

grecipe-manager.sourceforge.net

If you're chasing recipes from around the world and want a neat program with clever features to catalogue them all, Gourmet Recipe Manager (or Gourmet for short) is definitely your tool. According to its Freshmeat entry:

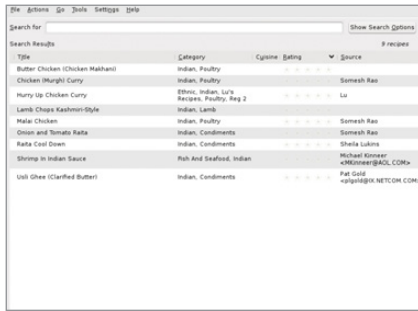
Gourmet Recipe Manager is a simple but powerful recipe-managing application. Gourmet is intended for the GNOME desktop environment, but works on any platform that GTK supports, including Windows.

Gourmet allows you to collect, search and organise your recipes, and to generate shopping lists automatically from your collection. The latest version also lets you calculate nutritional information for your recipes using the USDA food database (or entering custom information by hand).

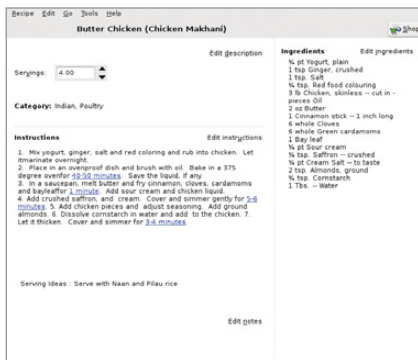
Installation Binaries are included for Debian- and RPM-based systems, as well as a source code tarball. The Web site states the following dependencies:

1. Python 2.3 (2.4 is recommended).
2. PyGTK > 2.3.9 and PyGnome. Note: with PyGTK > 2.5, PyGnome bindings have been split up. To print, you'll need, for example, the `python-gnome2-extras` package as well as the `python-gnome2` package.
3. libglade.
4. SQLite and Python SQLite2 bindings.
5. `distutils` (to install), provided by the `python-dev` package.
6. Python Imaging Libraries (PIL).

You'll also need the PyRTF library and the Gnomeprint bindings, but these are non-essential. I went with



By using openly available formats that are in use all around the world, tens of thousands of recipes are awaiting your download.



As well as being able to change the number of ingredients dynamically by altering the serving size, there are links to timing mechanisms all through the recipes.

the Debian package and found I needed to install python-sqlalchemy to get the package running on my Ubuntu system. If you go with the source version, download the tarball, extract it, and open a terminal in the new folder.

As root or sudo, enter the following:

```
# python setup.py install
```

Once that's out of the way, you now can run the program by looking in your system's menu (it was in Applications→Utilities→Gourmet Recipe Manager under KDE 4), or by entering the command:

```
$ gourmet
```

Usage Before you can start looking at recipes, you need to download some first. I had to search for a file and then import it before I could use anything. I did a search for "mastercook" recipe files (Gourmet also supports Meal-Master recipes, but I found the Mastercook recipes easier to deal

There's also a seriously cool function called **Servings**, which changes the quantities of the recipe's ingredients when you're serving more or fewer people than the recipe's default.

with), and found a Web page full of Mastercook recipes zipped up. After downloading the zip and extracting it, I imported the *.mxc file by clicking File→Import file.

Once you have imported a file successfully, there will be a new list of recipes, and you can double-click on any recipe to view it. As well as a recipe name, Gourmet helpfully includes food categories, the author of the recipe, and some even carry a star rating. Once you have opened a recipe, the instructions are presented neatly for you on the bottom left, with the ingredients (and quantity thereof) listed on the right.

Cool features I found include a unit

converter, the ability to generate a printable shopping list based on the recipe's ingredients, as well as a built-in timer. On the subject of timing, within each part of the recipe that involves performing a task over a period of time, there's usually a link to a timer that will be preset to however long you need to cook and so on. On each of these timers is a start, pause and reset button, allowing you to get the task timed and underway as you like.

There's also a seriously cool function called **Servings**, which changes the quantities of the recipe's ingredients when you're serving more or fewer people than the recipe's default. Simply

DEDICATED SERVERS WITH BLAZING FAST CONNECTIONS

Gigabit ports / MULTI-Gig options
 High-capacity bandwidth plans, including:

- * 3000 GB/month for \$200
- * 5000 GB/month for \$375
- * 10000 GB/month for \$800

Custom clusters with private VLANs
 Flexible storage and RAID options

Intel Premium Partner

Numerous OS choices (Linux or Windows)
 FREE 24x7 "6-Star" support

www.CARI.NET/LJ
888.221.5902

carinet
 Better Servers. Better Service

change the number of servings from the default number, and the ingredients list updates dynamically to tell you how much you need of each ingredient for the new number of servings.

Gourmet Recipe Manager is a great tool that is well thought out, and hopefully, it will find an audience among home enthusiasts and gourmet chefs alike. Given that it runs on any platform that supports GTK (especially Windows), hopefully its appeal and usage will extend far beyond just the Linux platform to home PCs everywhere. With the added bonus that there are tens of thousands of possibly supported recipes from all over the world, food in some homes and restaurants may just become a little more exotic.

gWaei—Japanese-English Dictionary

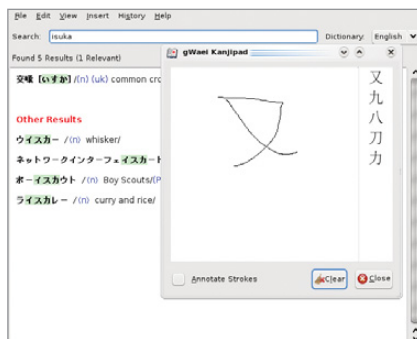
gwaei.sourceforge.net

Students of Japanese have had a number of tools available for Linux for sometime, but here's a project that updates the situation and brings several elements together from other projects to form one sleek application. In the words of the gWaei Web site:

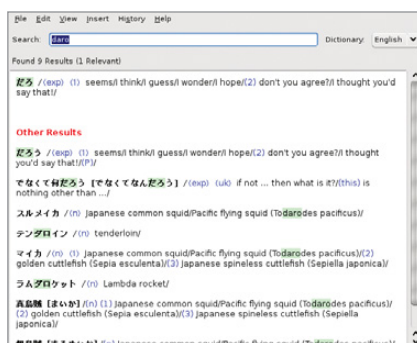
gWaei is a Japanese-English dictionary program for the GNOME desktop. It is made to be a modern drop-in replacement for Gjiten with many of the same features. The dictionary files it uses are from Jim Breen's WWWJDIC Project and are installed separately through the program.

It features the following:

- Easy dictionary installation with a click of a button.
- Support for searching using regular expressions.
- Streams results so the interface is never frozen.
- Click Kanji in the results pane to look at information on it.
- Simple interface that makes sense.
- Intelligent design and Tab switches dictionaries.



gWaei is a quick and easy-to-use Japanese dictionary that updates some old Linux favorites.



The coolest feature in gWaei is this kanji pad, where you can draw kanji with your mouse, and the computer dynamically alters the selection based on your strokes.

- Organizes relevant matches to the top of the results.

Installation If you head to the Web site's download section, there are gWaei packages available in deb, RPM and source tarball format. For me, the deb installed with no problems, so I ran with that. When running with the source version, I couldn't find all of the dependencies, but the Web site says you need the following packages, along with their -dev counterparts: gtk+-2.0, gconf-2.0, libcurl, libgnome-2.0 and libsexy.

The documentation also says that compiling the source is the standard fare of:

```
$ ./configure
$ make
$ sudo make install
```

After installation, I found gWaei in my menu under Applications→Utilities→gWaei Japanese-English Dictionary. If you can't find gWaei in

your menu, enter the command:

```
$ gwaei
```

Usage Once gWaei starts, the first thing you see is a Settings window that's broken into three tabs: Status, Install Dictionaries and Advanced. Status tells you how things are currently set up, and to start off with, all you'll see is Disabled. Click the Install Dictionaries tab, and you'll see that there are buttons already set up to install new dictionaries, called Add, for English, Kanji, Names and Radicals. Once these are all installed, each of them will be changed to Enabled back in the Status tab.

After these are installed, click Close, and you are in the program. The first place you should go is the search bar. Enter something in English or in Romaji (Japanese with the Latin alphabet we use), and meanings and translations appear in the large field below with a probable mix of kanji and kana, and an English translation. You also can enter searches in kana and kanji, but my brother has my Japanese keyboard, so I couldn't really try it out.

For a really cool feature, click Insert→Using Kanjipad, and a blank page comes up where you can draw kanji characters by hand with your mouse. Various kanji characters then appear on the right and update, depending on how many strokes you make and their shape. If you click Insert→Using Radical Search Tool, you can search for radicals on basic kanji characters, which also can be restricted by the number of strokes.

All in all, gWaei is a great program with elegant simplicity, and it has the features you need, whether you're in Japan or the West (or anywhere else that's not Japan for that matter). If you're a Japanese student, this should be standard issue in your arsenal. ■

John Knight is a 24-year-old, drumming- and climbing-obsessed maniac from the world's most isolated city—Perth, Western Australia. He can usually be found either buried in an Audacity screen or thrashing a kick-drum beyond recognition.

Brewing something fresh, innovative or mind-bending? Send e-mail to newprojects@linuxjournal.com.

Celebrating 15 years of *Linux Journal*,
we've brought together every article ever
published in the world's #1 Linux magazine
and packaged it in one
convenient, fully searchable CD.



With nearly 4,000 articles written by industry experts on everything from cool projects, desktop how-tos, security, embedded systems, networking, virtualization, multimedia, system administration and programming tricks and techniques—this unique collection is a must-have for every Linux enthusiast.

Get your **NEW *Linux Journal* Archive CD** today featuring
all issues from 1994 through 2008. **Just \$34.95.**

www.linuxjournal.com/archivecd

2009 Readers'
Choice



LINUX
JOURNAL

Readers' Choice Awards 2009

If you pick it, they will come.

The *Linux Journal* Readers' Choice Awards have become an annual ritual, almost as fun as the holiday season. Our editorial team members can't wait to get their hands on the results to see what products and tools from the Linux space are keeping you productive, satisfied and wowed. And, who better to ask than our readers, the most talented, informed and (nearly always for the better) opinionated group of Linux experts anywhere? These characteristics are what make the awards such a great snapshot of what's hot and what's not in Linux.

Before diving into the results, let me explain that the results, although insightful, inherently fail to capture the true diversity of preferences that exist in our community. I wish we had space to show you the chaotic yet amazing mish-mash of responses to each question. As we try to lasso you into organized responses, you are curiously galloping off in pursuit of your category-busting solution

as we hang on to the rope for dear life. The reality is that you are always experimenting; your opinions are fluid, and filling in virtual bubbles doesn't fully explain the nuance of your relationship to your tools.

This is a survey of big trends, and the trend underlying them all is that you embrace lots of tools. One respondent summed it up with "All of the above, many of the above", and another exclaimed, "Variety is the spice, baby!"

Once again, in this year's competition, we designated only one winner per category, with strong contenders receiving Honorable Mention awards. For instance, in the categories where a cluster of formidable contenders followed the outright winner, we designated up to three honorable mentions. Besides a few exceptions, a product or service had to get at least 10% of the vote for Honorable Mention status.

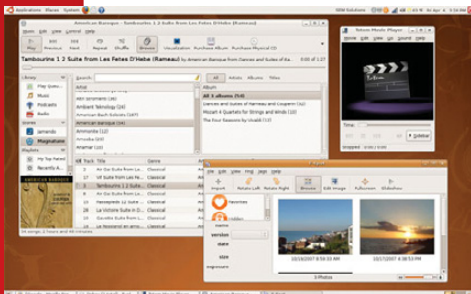
Here then, ladies and gentlemen, are *your* 2009 *Linux Journal* Readers' Choice Awards.

JAMES GRAY

Favorite Primary Linux Distribution of Choice

- Ubuntu** (45%)
Honorable Mention
- Debian** (10%)

The Ubuntu phenomenon, even by Linux standards, is truly paradigm-shifting and it doesn't appear that it will be tanking any time soon. In the 2008 Readers' Choice Awards, we dubbed the ubiquitous Ubuntu "the big distro that did" for unexpectedly leaving its myriad rivals as mere dots in the rear-view mirror. In the 2009 tally, Ubuntu deserves an "Energizer Bunny Award" for winning the Favorite Primary Linux Distro category, increasing its popularity and becoming untouchable for the time being. Ubuntu and its related Kubuntu and Edubuntu siblings raised their vote tally from 37% in last year's awards to 45% this year. At whose expense is Ubuntu gaining? Although Mandriva tumbled most dramatically from 14% to 2%, Fedora and Novell/SUSE hemorrhaged more than a few percentage points. Meanwhile, Red Hat, CentOS and Debian ticked up a few points from last year, the latter enough to warrant Honorable Mention with its 10% share of your vote.



Favorite Desktop Environment

- GNOME** (53%)
Honorable Mention
- KDE** (30%)

During the past year, GNOME has reached majority rule status, with 53% of you electing it your favorite desktop environment. This trend is despite the breakneck development of KDE 4 during the past year. Although GNOME garnered only a few more votes than it did in 2008, KDE's vote count slipped as you've warmed to Xfce, Fluxbox and Enlightenment. The long and influential coattails of Ubuntu can only make any presidential candidate green with envy.



Favorite Web Browser

- Firefox** (87%)

Firefox takes first prize as both your Favorite Web Browser for 2009 and the most extreme "category crusher" of this competition. No other application (besides the competitor-less Apache) racked up a higher share of votes (87%) than Firefox, although OpenOffice.org wasn't far behind. Who can argue, as Firefox keeps getting faster and accumulates more useful extensions? This year's surprise gainer was the historically underappreciated Opera browser, which ratcheted up from 5% to 8% on the heels of its impressive 9.5 and 9.6 releases. The worthy Konqueror and the browsers based on the Firefox Gecko engine (for example, Flock and Epiphany) were left behind in the catchall "Other" category. How will this category look next year? Look for an inevitable battle royale if Google can deliver a polished Chrome for Linux in time for you to give it a test drive.



Favorite E-Mail Client

- Mozilla Thunderbird** (38%)
Honorable Mention
- Gmail Web Client** (31%)
- Evolution** (11%)



Congratulations to Mozilla Thunderbird for its third consecutive win in the Favorite E-Mail Client category with 38% of the vote. The "Holy how did that happen, Batman?" award, as well as Honorable Mention, go to the runner up, Gmail, which got 11% more of your votes than last year, while Thunderbird dropped 7%. Meanwhile, other non-Web-based clients, Evolution and KMail, are ever more the favorites of fewer. Is the trend toward Web-based clients inexorable, or can the non-Gmails of the world find a formula to stanch the decline?

Favorite Office Program

OpenOffice.org (85%)



No news here, gang. OpenOffice.org, bolstered no doubt by its recent 3.0 release, retains its crown as your Favorite Office Program with the same share of your vote as last year, 85%. The alternatives, AbiWord and KOffice, each similarly continue to be the favorites of merely 3% of you. As with Firefox, there seems to be some sort of correlation between a program's cross-platform characteristics (should any exist) and its category-crusher status. Honorable Mention for most definitive response in this category goes to "I hate all office programs".

Favorite Audio Tool

Audacity (73%)



In this year's competition, we decided it made sense to split up audio-related programs into two categories. The first is Favorite Audio Tool—that is, program for creating, manipulating and modifying audio streams. The second is Favorite Audio Player—that is, program for playing and organizing existing audio streams. The cross-platform Audacity sound recorder and editor is yet another—cross-platform rule holds—category crusher in the Favorite Audio Tool department, garnering top marks from 73% of you. Although the applications LMMS and Ardour each have a critical mass of adherents, each is the favorite audio tool of only 6% of you.


Favorite Media Player

VLC (34%)

Honorable Mention

MPlayer (33%)

Ladies and gentlemen, we have a new champion in the Favorite Media Player category. VLC, which last year landed in a close second place behind MPlayer, won the category in a photo finish by a single percentage point. Both players play pretty much any format you can throw at them, making usability a key factor in your decision. Thus, more of you are opting for the single-console approach of VLC rather than MPlayer. Other players that recorded respectable results were Totem (9%), Kaffeine (8%), SMPlayer (6%) and xine (5%).



Favorite Audio Player

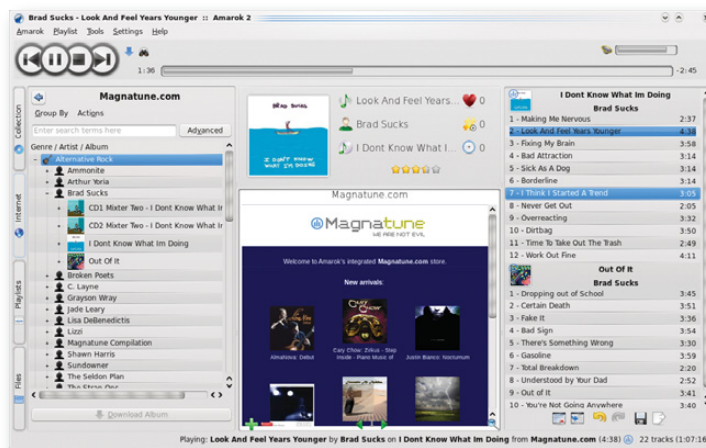
Amarok (36%)

Honorable Mention

Rhythmbox (18%)

XMMS (12%)

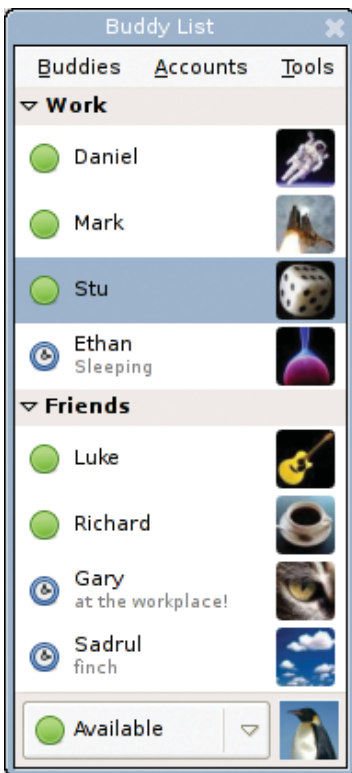
Although alternatives have knocked Amarok back a few points from last year, its 36% share of the vote tally helped the renowned audio player for KDE remain undisputed champion in the Favorite Audio Player category. Many of you also are aficionados of the audio players Rhythmbox and XMMS, each of which received enough votes to warrant Honorable Mention. It's interesting to see KDE flagging as your favorite desktop, yet the KDE audio player, admittedly more feature-packed, has double the support of the GNOME audio player, Rhythmbox.



Favorite Communications Tool

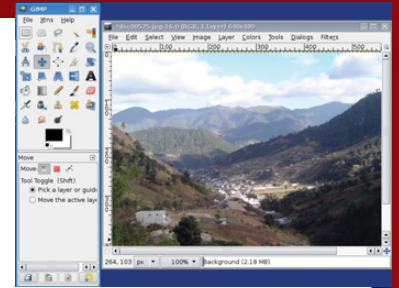
- Pidgin** (43%)
Honorable Mention
- Skype** (18%)

Nothing changed in the Favorite Communications Tool category. Once again, Pidgin Internet Messenger, the Swiss Army knife-esque messaging tool formerly known as Gaim, took top honors with a 43% share of your votes. Pidgin users appreciate the ability to monitor all of their messaging accounts in 15 different protocols. Competitor Kopete, which slipped from its Honorable Mention status from last year, does only 11 protocols. Meanwhile, the closed-source Skype retained its Honorable Mention laurels by earning 18% of your votes. That may be its ceiling until its improbable open-source resurrection, as a sizable contingent of us will never fill the bubble of any closed-source application, regardless of how good it is.



Favorite Graphics/Design Tool

- GIMP** (76%)
Honorable Mention
- Inkscape** (11%)



The legendary GIMP remains your unrivaled choice for Favorite Graphics/Design Tool, once again with 76% of the votes. It appears though, that the vector-graphics application Inkscape is emerging from the pack of graphics applications as a new favorite. Inkscape left the single-digit vote-getters to reach 11% of your votes, enough to win it Honorable Mention in the category. Are the impressive, mature programs like Blender too specialized to warrant your vote? Maybe the category is too broad. One write-in voter exclaimed, "You're making me choose between GIMP and Blender?!", and another explained, "Blender, GIMP and Inkscape are totally different tools for different purposes. They're all my favorites in their respective categories". Points well taken.

Favorite Digital Photo Management Tool

- Picasa** (34%)
Honorable Mention
- F-Spot** (17%)
- digiKam** (13%)
- gThumb** (11%)



Although much has changed in the crowded category of Favorite Digital Photo Management Tool, one broad trend appears to hold. If a Google application is around, it is likely to be slicing and dicing its rivals. In the photo management category, the slicer-dicer is Picasa, and the sliced and diced is digiKam. In the 2008 awards, Picasa and digiKam were neck and neck with 25% of the votes. This year, Picasa wins the category, leaving everyone else in the dust with its 34% of the vote. F-Spot (at 17%), digiKam (at 13%) and gThumb (at 11%) are still all popular enough to deserve Honorable Mention. However, Picasa may continue to surge as unique features, such as the ability to sync photos between one's PC and Web-based albums seamlessly, make it a tough act to follow.

Favorite Text Editor

- vi** (36%)
Honorable Mention
- gedit** (19%)
- Kate** (11%)

The more things change, the more they stay the same in the Favorite Text Editor category. Vi wins again with a solid 36%, with gedit and Kate taking Honorable Mention honors. Emacs and nano also are popular but just missed the cut.

Favorite Version Control System

Subversion (47%)

Honorable Mention

CVS (16%)

git (15%)



You left little doubt about who deserves to win Favorite Version Control System, a new category in the 2009 Readers' Choice Awards. Subversion is the favorite of 47% of you; CVS and git win Honorable Mention at 16% and 15%, respectively.

Favorite Database

MySQL (61%)

Honorable Mention

PostgreSQL (18%)

MySQL's move over to Sun Microsystems doesn't appear to have dampened your admiration for the legendary open-source database. Both this year and last year, you deemed MySQL your Favorite Database, with 61% of your votes this year. PostgreSQL also shared a similar fate as last year, registering 18%, enough for Honorable Mention. SQLite, Oracle and Firebird all polled in the single digits.

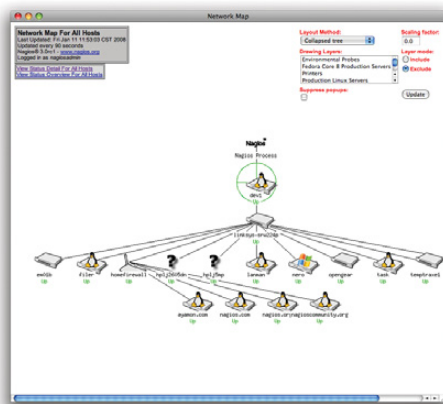
Favorite Linux Monitoring Application

Nagios (51%)

Honorable Mention

Hyperic HQ (15%)

up.time (11%)



Nagios was not only recently dubbed one of the most important open-source apps of all time, but it also is the winner of the new Readers' Choice category, Favorite Linux Monitoring Application. A slim majority 51% of you use Nagios to keep close tabs on your networks of all shapes, sizes and levels of complexity. Most of you not using Nagios opt for the Honorable Mention candidates, Hyperic HQ (with 15%) and up.time (11%). Ganglia and GroundWork also garnered respectable votes in the single digits.

Favorite Programming Language

Python (20%)

Honorable Mention

C++ (19%)

Java (17%)

C (13%)

Perl (12%)

Last year, we created discord when we split programming languages into two categories: Favorite Programming Language and Favorite Scripting Language. Then, we limited your choices according to our own definition of each. In order to shield ourselves from the avalanche of "WTFs" (whew, we succeeded!), we gave you more latitude to decide which is which. Therefore, the results look a bit different from last year. In an interesting twist, Guido van Rossum's venerable Python, which took First Place in last year's Favorite Scripting Language category, wins this year's Favorite Programming Language award with a hefty 20% of your votes. Close behind in the Honorable Mention group are your other favorites, with few surprises: C++ with 19%, Java with 17%, C with 13% and Perl with 12%.



Favorite Scripting Language

- bash** (28%)
Honorable Mention
- PHP** (24%)
- Python** (19%)
- Perl** (14%)

The results of the Favorite Scripting Language illustrate the diversity of opinions on what is a scripting language. Although the prosaic workhorse bash (shell) wins the category with 28% of the tally, three other quite different languages follow close behind in the Honorable Mention category: the Web-centric PHP, the flexible Python and the Swiss Army chainsaw of programming languages, Perl.

Favorite GUI Remote Access or Network Computing Solution

- SSH and X** (40%)
Honorable Mention
- TightVNC** (14%)
- rdesktop** (13%)
- RealVNC** (12%)

Your inaugural choice for Favorite GUI Remote Access or Network Computing Solution is clear. SSH and X wins hands-down with a commanding 40% share of the votes. Meanwhile, a hefty chunk of you choose to go graphical, using variants of VNC, such as TightVNC, RealVNC and UltraVNC. In fact, if you add those three user groups together, you're just shy of winning the category. TightVNC, rdesktop and RealVNC are all popular enough to share the platform for Honorable Mention.

Favorite Linux IDE

- Eclipse** (42%)
Honorable Mention
- NetBeans** (14%)
- KDevelop** (11%)



Yet another new category in this year's awards is Favorite Linux IDE, which the ubiquitous Eclipse won commandingly and unsurprisingly with 42% of the votes cast. The fact that in Eclipse one can work in a lean environment and add and subtract an incredible array functionality with its myriad modules has closed the deal for nearly a majority of you. At the same time, the second largest vote-getter was "Other". Clearly the Linux developer community cannot be pigeonholed.

Favorite Platform for Developing Rich Internet Apps

- Adobe Air** (21%)
Honorable Mention
- Gears** (18%)
- JavaFX** (15%)



When it comes to your Favorite Platform for Developing Rich Internet Apps (yet another new category for 2009), you are less decided than in the Linux IDE category. Although Adobe Air is the favorite of the most of you at 21%, you also are using Gears and JavaFX in solid numbers, 18% and 15%, respectively, among others. Mono Moonlight and OpenLaszlo also were close to the 10% mark. Will one of these tools break away to be the next Eclipse in a few years? Tune in to this space next year to find out.

Favorite Linux Game

- Frozen Bubble** (17%)
Honorable Mention
- Doom** (11%)
- Tux Racer**, also **Planet Penguin Racer** and **Extreme Tux Racer** (10%)



With some barely perceptible percentage changes, the Favorite Linux Game category remains the same as last year, led by *Frozen Bubble* and with Honorable Mention going to *Doom* and the *Tux Racer* series. Besides being consistent, the Favorite Game category is characterized by having the largest share of "Other" votes, with 27%, and the wittiest comments. One of you commented "Keeping it old school with SCUMM[VM] games". On the flip side, a surprising number of you also commented that you "have no time for games" or "don't like games". Meanwhile, this writer is wondering whether the many commercial game companies that now make Linux versions will ever break through with a runaway hit that could give *Frozen Bubble* a challenge one day.

Favorite Virtualization Solution

VirtualBox (32%)

Honorable Mention

VMware (30%)

Wine (13%)

We've been watching VirtualBox for a few years now, wondering when its popularity would finally match its technical prowess. Well, 2009 is finally VirtualBox's time in the sun, as this year it toppled VMware to win Favorite Virtualization Solution. Last year, VirtualBox received roughly half the votes of VMware (20% vs. 39%). This year, VirtualBox won the matchup 32% to 30%. VMware and Wine, thus, took Honorable Mention honors. Xen fell just short of 10% of the vote. [See Marcel Gagné's *Cooking with Linux* on page 26 for more on VirtualBox.]



Favorite Package Management Application

apt (37%)

Honorable Mention

Synaptic (16%)

Yum (13%)

RPM (10%)

As the fate of Ubuntu and siblings slopes ever upward, so too trends the popularity of apt, the principal package management system for Ubuntu and Debian. Once again, apt, with 37% of the votes, is the clear and even more dominant victor as Favorite Package Management Application. More of you also are turning to Synaptic, the groovy front end to apt, to keep your system loaded with your favorite programs. The classic RPM (10%) and its amigo Yum (13%) declined slightly from last year, as more of you are leaving Fedora and Novell/SUSE in favor of Ubuntu.



Favorite Backup System

Amanda (16%)

Honorable Mention

Bacula (14%)

Simple Linux Backup (14%)

Once again, in the backup department we differentiated between comprehensive applications, or systems (this category), and specific utilities (see Favorite Backup Utility below). Regarding the Favorite Backup System category, most of you continue to prefer the same systems as last year, although the deck chairs have been rearranged a bit. This year, the open-source application Amanda took the victory lap, which Simple Linux Backup took last year. Still, the latter won Honorable Mention this year along with Bacula (yet again) for network-based backup. Favorite Backup System also has a variety of write-in votes with some variation of "roll my own solution".

Favorite Backup Utility

rsync (47%)

Honorable Mention

tar (34%)

Though rsync and tar are your perennial favorites for Favorite Backup Utility, this year, the two flip-flopped positions, with tar taking the crown last year and rsync taking it this year. rsync is the favorite backup utility of 47% of you to tar's 34%.

Favorite Content Management System

WordPress (25%)

Honorable Mention

Joomla! (23%)

Drupal (19%)



No surprise that the blog publishing application WordPress once again wins the category Favorite Content Management System with 25% of the vote. In the Honorable Mention department, your same two favorite Web content managers, Joomla! and Drupal, are present only to flip in popularity. This year, Joomla! reached an impressive 23% to Drupal's 19%. Drupal got our vote—LinuxJournal.com runs on it.

Favorite Web Server

Apache (89%)

As with last year, the thought arises in Favorite Web Server category of whether we should just ask "Do you use the Apache Web server, yes or no?" and leave it at that. Apache wins again in 2009 with 89% of your votes.



Favorite Linux-Friendly Web Hosting Company

Contegix (15%)

Honorable Mention

Rackspace (12%)

Talk about a meteoric rise, Contegix went from one write-in vote in 2008 to champion of the Favorite Linux-Friendly Web Hosting Company category in 2009. Otherwise, four well-known names have remained among your perennial favorites for years, namely Rackspace, GoDaddy.com, DreamHost and 1&1. Of the four, only Rackspace broke the 10% barrier this year for Honorable Mention, while the other three were just shy of the mark. GoDaddy.com was category winner last year, and DreamHost and 1&1 were the two Honorable Mentions.




Favorite Linux-Based Gadget

ASUS Eee PC (24%)

Honorable Mention

Android G1 (22%)

Last year, we admittedly were dorks for having the category Favorite Linux Handheld Device, which left so many cool Linux gadgets out in the cold. The Nokia N800 won that one. This year, however, we've taken our smart pills and expanded the category, calling it Favorite Linux-Based Gadget. Not surprisingly, one of the most well-known devices, the ASUS Eee PC won the category with 24% of the vote, followed by the Android G1, which achieved Honorable Mention status with 22%. After the G1, the field was so crowded, making it impossible, unfortunately, for any other device to crack the 10% barrier. The Nokia N810 Tablet, Acer Aspire One, TomTom Navigation System, OpenMoko FreeRunner, Amazon Kindle, the Palm Pre and several write-ins all received a fair share of your vote, which shows how sophisticated, interesting and crowded the Linux device space has become.



Favorite Linux Laptop

ASUS Eee PC (32%)

Honorable Mention

Lenovo T61p (16%)

Dell Inspiron Mini 9 (12%)

Acer Aspire One (10%)

There is something oddly liberating about the "big guys" pre-installing Linux on their PCs. At long last, when we go to buy a PC, a device so central to our identities and livelihoods, we find the well-thought-out preference for Linux taken seriously by the companies we want to buy from. After being shut out so long for being too smart, it sure feels good, doesn't it? Your vote for the ASUS Eee PC as Favorite Linux Laptop (with 32% of the vote) tells us how much you appreciate the opportunity to buy a laptop designed with Linux in mind and not just a feature-handicapped afterthought to placate the pesky geeks. The group of Honorable Mentions includes not only the returning Lenovo T61p (16%) but also the newcomers Dell Inspiron Mini 9 (12%) and Acer Aspire One (10%). Despite such euphoria over the big guys, you didn't forget our Linux-specialist friends like Linux Certified, EmperorLinux and R Cubed who kept us motoring during darker times. They fared well as a group if you add up all the votes for their various models.

Favorite Linux Desktop Workstation

Dell (41%)

Honorable Mention

Hewlett-Packard (16%)



Let's start with the official results for Favorite Linux Desktop Workstation. Dell won the category with 41% of your votes, and Hewlett-Packard earned Honorable Mention with 16%. Unfortunately, the official results fail to appreciate the "roll-your-own" spirit that is so vital to our community. Because we didn't include a choice like "I configure my own desktop PCs", you told us as much in your own words. In a classic survey creator's nightmare, the responses "I do. :)" and "I do, as in self-built" and "Home-brewed" all registered as separate votes worth 0.05% each even though they mean the same thing. Allow me put on my Katherine Harris hat and have a look at these "hanging chads", Florida-election-style, to shed some light on your roll-your-own tendencies. Hours of investigative sleuthing revealed that roughly 12% of you configure your own desktop PCs. Therefore, the honorary Honorable Mention award in this category goes to the roll-your-own spirit of the Linux Community.



Favorite "Green" Linux Product or Solution

Virtualization (45%)

Honorable Mention

PowerTOP Tool (16%)

Last year, VMware took top honors as Favorite "Green" Linux Product or Solution partly because of how we phrased the question. This year, to be more fair, we grouped virtualization solutions together, and they won the category with 45% of your votes. The win makes sense given the technology's impressive improvement in the efficiency of servers. The PowerTOP tool for finding energy wasters on your systems also is popular and won Honorable Mention with 16%. We failed to list the recent (kernel 2.6.21) innovation of the tickless idle on Linux, which takes advantage of low power states in modern processors. Are you taking advantage of this feature? Next year, we'll ask you directly. Finally, this author wishes to express his dismay at the significant number of disparaging remarks in this survey toward green solutions. Although the vast majority of respondents are positive to neutral in this category, responses such as "Don't drink the green Kool-Aid" and "I don't care!" were plentiful. Will our progeny admire our arrogant proclivity to waste natural resources and do little to change our ways? I doubt it.



Favorite Linux Server

Dell (32%)

Honorable Mention

IBM (16%)

Hewlett-Packard (15%)

While the roll-your-own philosophy is alive and well when it comes to servers, you tend to feel more comfortable giving this business to the big guys. Dell is the winner of the Favorite Linux Server category with 32% of your votes. Your Honorable Mention winners, IBM and HP, trailed Dell with 16% and 15%, respectively.

Favorite Linux Book of All Time

Linux in a Nutshell by Ellen Siever, Stephen Spainhour, Stephen Figgins and Jessica P. Hekman (4%)

Honorable Mention


Just for Fun: The Story of an Accidental Revolutionary by Linus Torvalds and David Diamond (3%)

Running Linux by Matt Welsh, Matthias Kalle Dalheimer, Terry Dawson and Lar Kaufman (3%)

The Cathedral and the Bazaar by Eric S. Raymond (2%)

The Linux Bible by Christopher Negus (2%)

Last year, we limited the selection in this category to books published from 2007 to the date of the survey (February 2008). This year, we decided to try Favorite Linux Book of All Time, and the results didn't gel as easily. Because this category is so crowded, we decided to toss out the 10% rule for Honorable Mention and honor your top five books. Interestingly, two of your favorites were not technical guides but rather required reading for understanding the Linux phenomenon, namely *Just for Fun: The Story of an Accidental Revolutionary*, the story of Linus Torvalds' rise to fame, and *The Cathedral and the Bazaar*, Eric Raymond's seminal book on the Open Source movement. Rounding out the Honorable Mentions *Running Linux* and *The Linux Bible*. The information-overloaded among you meanwhile opted for the popular write-in candidates "Too many to choose" or "Very tough question".



Favorite Linux Journal Column

Tie: **Marcel Gagné's Cooking with Linux** (19%) and **Kyle Rankin's Hack and /** (19%)

Honorable Mention

Dave Taylor's Work the Shell (10%)

Mick Bauer's Paranoid Penguin (10%)



Here is some *Linux Journal* lore for you. Our publication used to have a column called Kernel Korner, which many faithful readers certainly remember. Kernel Korner dominated the Favorite *Linux Journal* Column category from the awards' inception in 1996 until 2000. Things got competitive in 2001 when Marcel Gagné and his *Cooking with Linux*, the world's first column ever to pair practical (and hilarious) Linux advice with appropriate wine selections. *Cooking with Linux* has been so popular ever since that it won Favorite *Linux Journal* Column from 2001 to 2008, and last year, I suggested that "Marcel Gagné is going to have to be knocked off before anyone knocks him off the award stand." My prediction was premature, because this year, although Marcel won this category again, he shares his title with Kyle Rankin's more recent *Hack and /* column. Believe it or not, Marcel and Kyle received the exact same number of votes, or 19% each. Congratulations are due to both excellent columnists who offer vastly different but equally useful content in their monthly musings. Meanwhile, Dave Taylor's *Work the Shell* and Mick Bauer's *Paranoid Penguin* are both popular with 10% of you, enough to award them Honorable Mention.

LINUX PRODUCT OF THE YEAR

Android Platform and the T-Mobile G1 Phone (9%)

Honorable Mention

Ubuntu (8%)

KDE 4 (7%)

ASUS Eee PC (6%)



In the question for 2009 Linux Product of the Year, we didn't give you any suggestions. We left the responses 100% up to you. Naturally, this made nearly every response unique and left it up to us to categorize it. Nevertheless, it is safe to proclaim that your 2009 *Linux Journal* Product of the Year Award goes to the Android platform and its first commercial implementation, the T-Mobile G1 phone. The pair garnered 9% of your votes. Close behind, with 8% and Honorable Mention, was the Ubuntu Linux distribution, followed by the KDE 4 desktop with 7% and, finally, last year's winner, the ASUS Eee PC with 6%. Last year, the Eee PC reached an impressive 37% of the votes. Interestingly, the development of Android and the G1 phone, although popular and groundbreaking, didn't have quite the same dominant effect that the Eee PC had last year—nor did any other single product. This effect allowed you to remediate my lament from last year that "when Ubuntu releases yet another fantastic upgrade, our expectations are met and the buzz meter quickly subsides". This year, Ubuntu got the respect it deserves in this category for revolutionizing the Linux desktop oh so gradually with each great upgrade.

Thanks to each and every one of you who participated in the voting. ■

James Gray is *Linux Journal* Products Editor and a graduate student in environmental sciences and management at Michigan State University. A Linux enthusiast since the mid-1990s, he currently resides in Lansing, Michigan, with his wife and cats.



“Accomplishing great things with limited resources is the open source way of life.”

—Allison Randal,
OSCON Program Chair

Register now for **OSCON 2009**, save \$250, see San Jose, and discover why open source is open for business.

OSCON 2009 will tackle the hard questions and propose some answers: How can open source — its tools and its principles — contribute to making a difference, both in the business of computing and in creating a sustainable lifestyle? In an uncertain economy, how can open source empower us? How can we turn up the volume on efficiency, knowledge transfer, and working smarter within constraints to achieve more with what we already have — or even with less?

Join us as we redefine the future of open source and the savings it offers to all of us. **For starters, you'll save \$250 when you register for OSCON by June 2, at:**
<https://en.oreilly.com/oscon2009/public/register>



Register by June 2, 2009 and **SAVE UP TO \$250**
Save an additional 15% off registration by using code OS09LJ

Building a Linux-Based High-Performance Compute Cluster

The Rocks clustering package from the University of California at San Diego makes it easy to build and maintain a high-performance compute cluster with off-the-shelf hardware. TOM LEHMANN

You have an application running on a relatively new dual-core workstation. Unfortunately, management wants it either to complete faster or be able to take on a larger dataset in the same time as it runs now. You do a bit of investigating and find that both SMP and cluster versions of the application are available. You are using the SMP version on the workstation. You could speed things up if you could run on a quad core (or more) workstation, but the boss is not too receptive to the expenditure in the current economic climate. But wait, you do have a pile of 32 single-socket servers that were replaced earlier in the year. They're only single core, but 32 of them should have more capacity than the dual-core workstation, if you can just find a way to get them all to play together—that would be a cluster.

So, what is a cluster? Here's one

accepted definition: a cluster is a group of computers all working together on the same problem. To accomplish this, the machines in the cluster must be appropriately interconnected (a network) and trust each other.

It is possible to configure the networking and security manually, but there are easier ways to accomplish this using any one of a number of cluster provisioning and management systems. At the moment, one of the more popular packages is the Rocks package maintained by a team at the University of California, San Diego, under a grant from the National Science Foundation.

Rocks is termed a cluster provisioning, management and maintenance package. It helps you set up the cluster in the first place (from bare metal); it provides the tools to run parallel programs, and it provides the tools to maintain and extend

the cluster after it is created.

The package is delivered as a series of .iso images that you burn onto a series of CDs or DVDs. You then boot the machine that will become the head node from the appropriate DVD or CD, and the installation routine guides you from there. After asking a minimum number of questions in an interactive phase, the installation program builds the head node. Upon reboot, you invoke a single routine (insert-ethers) to add the rest of the machines as compute nodes. To add a compute node, you simply network boot it, and it will be added to the cluster, loaded and configured automatically. After the last node is complete, you have a functional cluster, ready to execute parallel applications.

So, with all of this in mind, let's build a cluster with those otherwise unloved machines.

STEP 1 Hardware Setup

The first item on the agenda is setting up the hardware. The overall idea is to have a set of connected computers. Ideally, the machines in the cluster should be as identical as possible, so no single machine or group of machines will be the weak link in any parallel computation. The same homogeneity should apply to the network, because most parallel computation relies on continuous communication between all of the nodes within the cluster.

Find a spot to set up your 32 servers, and make sure you have enough power and cooling to support them. As you connect all of the servers

to power, label both ends of each power cord so you can keep track of what is connected to each power strip in the rack.

Because you are starting with a clean sheet, now is a good time to update and configure the BIOS on each system. Set the BIOS clock to the current time as closely as practical (plus or minus five minutes is a good goal). Most clustering packages keep the BIOS clocks synchronized during operation, but only if the clock is reasonably close to the correct time at the beginning.

Because the machines are used, it's prudent to wipe all the disks before

loading the cluster software. There are many ways to accomplish this. One fairly thorough method is to use DBAN (Darik's Boot and Nuke). This self-contained application can perform several disk wipe techniques, including two that have some level of Department of Defense approval.

Remember, the goal here is to make all the machines in the cluster as identical as possible. But, this is a goal, not a hard and fast requirement. Heterogeneous clusters will work, but you may need to be careful as to how you deploy workloads on the machines to get the best performance.

STEP 2 The Network

Now that you have all the compute nodes configured and in the rack, it's time to set up the communications network. Figure 1 shows a typical networking setup for a simple compute cluster. In this configuration, the Ethernet fabric most likely would be used for administrative purposes, while the InfiniBand fabric would carry the compute traffic. If you don't have InfiniBand hardware available, you can

just ignore the bottom section of the diagram. The Ethernet fabric can carry both the administrative and compute traffic.

The best Ethernet network configuration for your cluster would be a single 48-port switch. If a switch like that is not available, you always can resort to a set of smaller federated switches forming a full fat tree network for the cluster. Like the compute nodes themselves, the network

should be as uniform as possible.

Plan all the cable runs, remembering that Ethernet cables have a nonzero cross section. Before you install them, test each cable. There is nothing as aggravating as finding that a cable is bad *after* it has been tied into the rack in a dozen places. Once again, label *both* ends of each cable to make troubleshooting simpler if it is necessary.

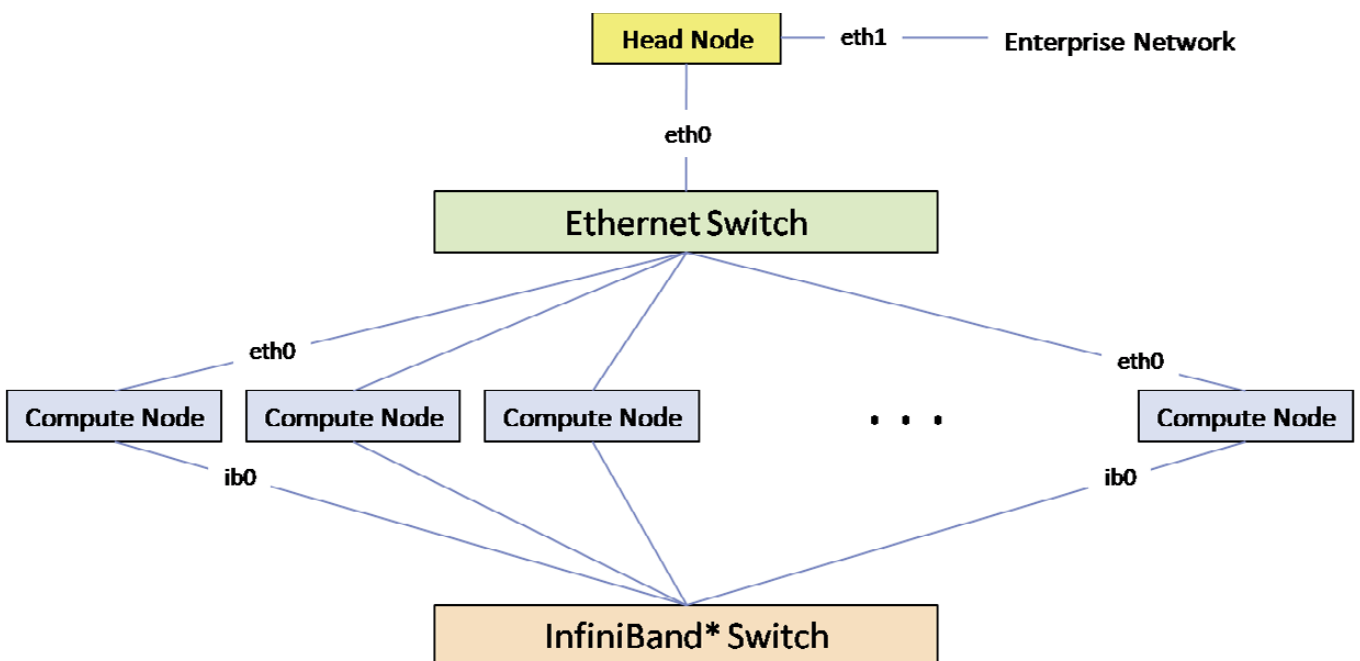


Figure 1. Network Setup for a Compute Cluster

STEP 3 Final Hardware Setup

Select one machine to be the head node for the cluster. The rest of the machines will be the compute nodes in your new cluster. As it installs the compute nodes, Rocks numbers the machines as compute-x-y, where the x is the rack number, and y is the number of the machine within a rack.

Say you spread the 32 nodes over four racks. If you want to follow the Rocks naming convention, you would set things up as follows: rack 0 contains the head node, so the numbering of the nodes would be compute-0-0 to compute-0-6. Rack 1 would come out as compute-1-0 to

compute-1-7. Rack 2 would contain compute-2-0 to compute-2-7. Rack 3 would follow suit.

Alternatively, you simply could pretend that all the machines are in a single rack: compute-0-0 to compute-0-30. Either way works, so use whatever is comfortable for you.

STEP 4 Get the Rocks Software Package

First, you need to get a copy of the Rocks package that will be appropriate for your cluster's hardware. Navigate to the Rocks Web site, and select the Download tab at the top of the home page to access the

various versions of the package. The 5.1 version is the latest at the time of this writing. Click the link to get a listing of the components of the package. For this exercise, I selected the x86-64 Jumbo DVD

image, downloaded it and burned it onto an empty DVD. While you are at the site, download the documentation. If nothing else, it will give you something to read while the software loads.

STEP 5 Boot Head Node and Select Installation Source

Boot the head node from the newly minted Rocks DVD. If everything is working as it should, you will be greeted with the welcome screen shown in Figure 2.



Figure 2. Rocks Welcome Screen

Enter `buil d` at the boot: prompt to start the installation sequence. The system boots in the normal Linux fashion and eventually presents the user with the initial Rocks configuration screen shown in Figure 3.

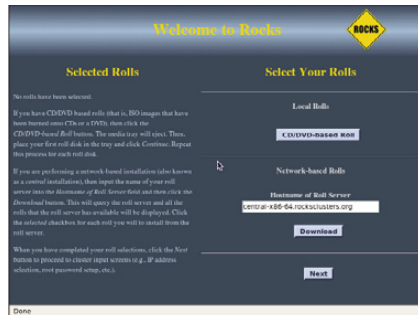


Figure 3. Rocks Configuration Screen

Because you have all of the software components on the Jumbo DVD, you will do your installation from CD/DVD-Based Rolls. Select CD/DVD-Based Roll. This brings up a screen listing all the individual components you can select from the DVD (Figure 4).



Figure 4. Rolls Selection Screen

For the purposes of this installation, I selected everything except the Bio Roll and the Virtualization Roll. You probably will select a different set of components. At absolute minimum, you need to select the Base, Web Server, Kernel and OS Rolls. Once you have made your selections, click Submit to continue the installation.

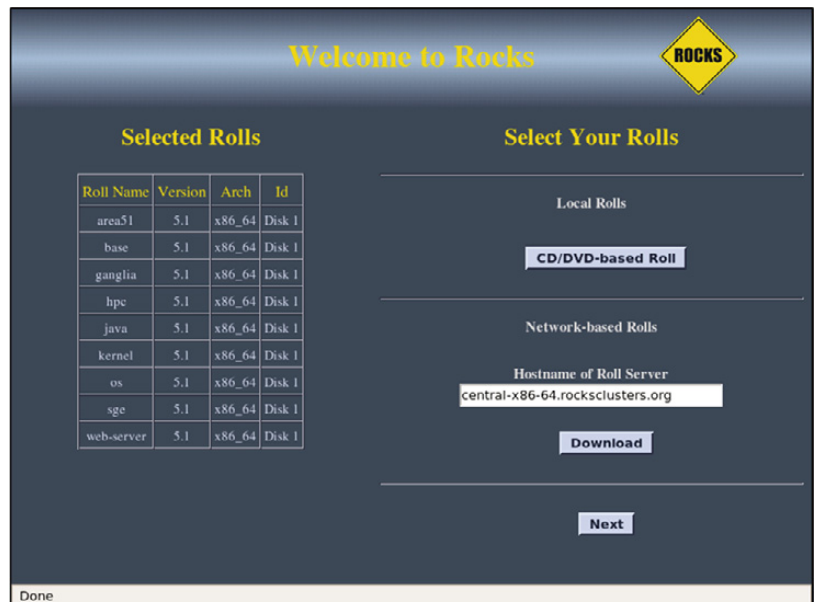

A Note about Rolls

The Rocks package is composed of a series of rolls. Some of the rolls are central to the cluster system itself (the Base Roll, the OS Roll, the Kernel Roll and the Web Server Roll). Some contain cluster functionality (the SGE Roll, the Java Roll, the HPC Roll and the Ganglia Roll). And finally, some rolls contain application software (the Bio Roll). Each roll is documented, so you can decide for yourself whether you need it in your installation. The commercial version of the Rocks package, Rocks+ from Clustercorp, adds additional packages that include commercial compilers from Absoft, Intel and the Portland Group as well as the TotalView debugger.

STEP 6 Confirm Your Selections

The installation now repeats the first screen, showing your selections (Figure 5). If you are satisfied with your selections, click Next to continue to the first of the administrative screens in the installation. If you want to make a change, click CD/DVD-Based Roll to go back to the component selection screen.

Figure 5. Rolls Selected Items Screen



STEP 7 Configure the Cluster

As you enter data on these screens, the installation routine is building a small MySQL database that details all of the component configurations in your cluster. The various tables Linux needs to run (like /etc/hosts) will be generated as an SQL report from this database. If you want to make changes in the system's configuration, the tools that Rocks provides actually change the database first, then run the appropriate reports to regenerate the system configuration files. This significantly reduces the chance for errors to creep into these files. It still is possible to edit the automatically generated system files manually, but remember that the next time you use the Rocks tools to reconfigure the cluster, your manual changes will be overwritten by the automatically generated SQL report versions.

The next screen (Figure 6) allows you to enter information about your cluster. If the cluster will be connected to your enterprise network, you should enter a fully qualified hostname to be consistent with your domain. The cluster name you enter in the Cluster Name field will appear in the management screens during cluster operation. Once you are satisfied with your entries, click Next to go to the configuration of the head node network connection to the private network (eth0).

Cluster Information	
Fully-Qualified Host Name	cluster.test.pro
Cluster Name	Cluster1
Certificate Organization	SDSC
Certificate Locality	San Diego
Certificate State	California
Certificate Country	US
Contact	admin@place.org
URL	http://www.place.org/
Latitude/Longitude	N32.87 W117.22

Figure 6. Cluster Information


Expert included.

Art is the Silicon Mechanics education and research expert. His mission is to consult with academic and research institutions and offer them the most compute power they can get for their money. Recently he's been talking with them about significant advances in personal supercomputing.

The Hyperform HPCg A2401 from Silicon Mechanics is a personal supercomputer with NVIDIA® Tesla™ GPU technology. This workstation starts with the AMD Phenom™ X4 processor, 8GB of DDR2 RAM, and it supports up to 8 hot-swap hard drives. With the addition of the NVIDIA Tesla C1060 GPU (or two, or three), the A2401 can outperform a small cluster—and it can do it without a cluster's noise, complexity, or cooling requirements. Best of all, it can do it without a cluster's price tag: the A2401 starts at a very user-friendly \$3139.

When you partner with Silicon Mechanics, you get more than high-end compute power at astonishingly affordable prices—you get an expert like Art.

For more information about the Hyperform HPCg A2401 visit www.siliconmechanics.com/TeslaPSC.

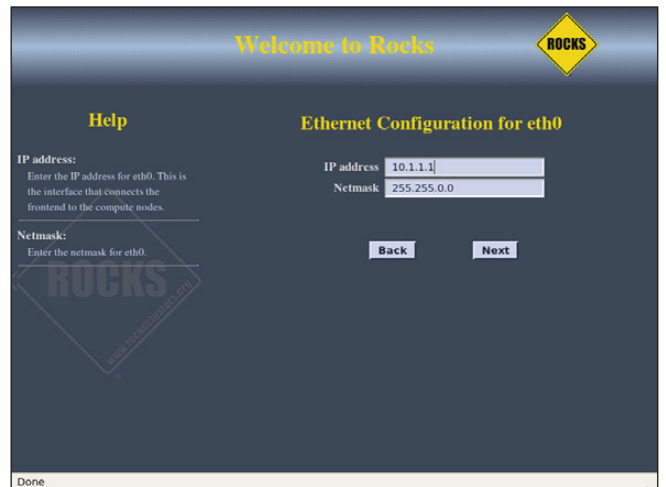


Silicon Mechanics and the Silicon Mechanics logo are registered trademarks of Silicon Mechanics, Inc. AMD, the AMD Arrow logo, AMD Phenom, and combinations thereof are trademarks of Advanced Micro Devices, Inc.

STEP 8 Configure the Cluster's Network

The next screen (Figure 7) lets you configure the cluster's network. The installation routine automatically selects 10.1.1.1 as the IP address for eth0 on the head node. Because this is a private network, you probably won't need to change this setting. If your public network also happens to be in the 10.1.X.X configuration, change this to something that doesn't conflict with your existing network. Clicking Next brings up the head node public network connection configuration screen.

Figure 7. Network Configuration



STEP 9 Configure the Public Network

Figure 8 shows configuring the "public" connection of the head node, its connection to the rest of your systems. The public connection for the head node must be configured with a fixed IP address. The public network for this example is configured as 192.168.0.X with a netmask of 255.255.255.0. Make sure the head node does not conflict with other servers and workstations on the public network. On the following screen (Figure 9), configure the local Gateway and DNS Server IP addresses for the head node to use.

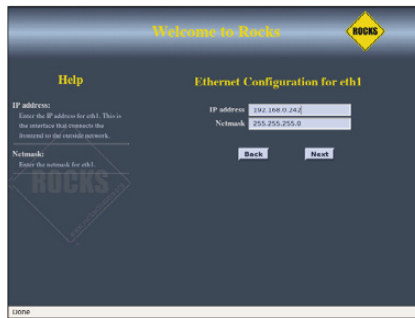


Figure 8. Head Node Public Network Configuration

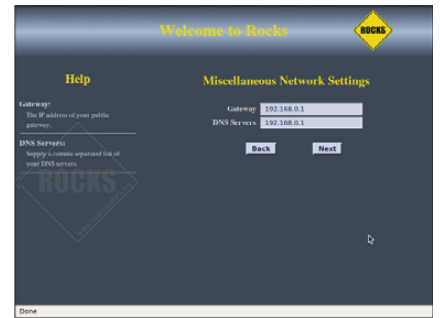


Figure 9. Head Node Gateway and DNS Configuration

STEP 10 Configure the Root Password and Time Zone

On the next two screens (Figures 10 and 11), enter the root password and configure the time zone and NTP server for the head node.

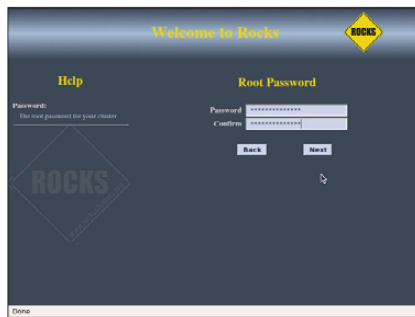


Figure 10. Root Password

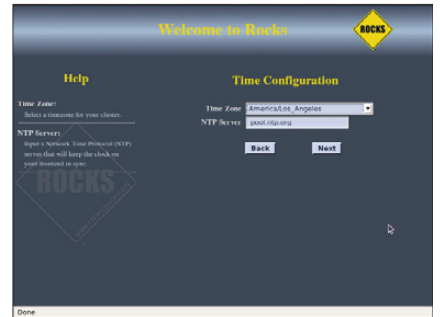


Figure 11. Time Zone and NTP Server

A Note about Time

All of the systems in the cluster must be synchronized as closely as possible to each other. This is accomplished using the Network Time Protocol (NTP). The head node synchronizes itself with one of the members of the public NTP pool then acts as a local time server through the private network connection to each member of the cluster. If a member of the cluster is slightly slow or fast, the NTP daemon on that machine will "slew" its clock over a period of time to bring it in line with the rest of the cluster.

STEP 11 Configure the Disk Partitioning

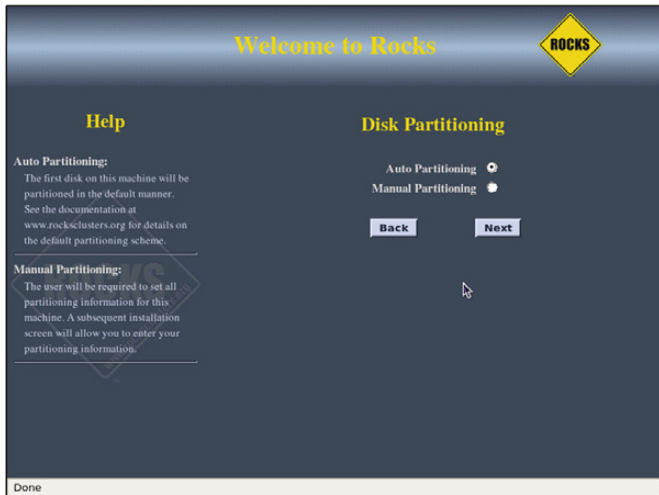


Figure 12. Disk Partitioning

The final interactive screen of the installation sequence (Figure 12) is the disk-partitioning screen. You can partition the disks automatically or manually. If you go with the automatic partitioning scheme, the installation routine sets up the first disk it discovers as follows:

Partition	Size
/	16GB
/var	4GB
swap	Equal to RAM size on the head node
/export (aka /state/partition1)	Rest of root disk

If you have multiple disks on the head node or you want to arrange the disk in a different fashion, select Manual Partitioning. This takes you to the standard Red Hat manual partitioning screen where you can configure things any way you desire (you still need to have a 16GB / partition and an /export partition at minimum though). Clicking Next on the disk-partitioning screen

Expert included.



As the head of Sales Engineering for Silicon Mechanics, Ken spends his time developing systems and configurations that are directly responsive to our customers' requests. That gives him unique insight into technologies that are catching on and gaining momentum. Lately Ken has been engineering a lot of clusters, and they tend to have some things in common. First, they are intended for use at a department or workgroup level. Second, they must be powerful but compact. Third, they need to be turnkey systems running Linux and the ROCKS+ cluster platform by Clustercorp Inc. Finally, they need to be reasonably priced.

Meet the Hyperform ROCKS+ Integrated cluster by Silicon Mechanics. Hyperform ROCKS+ Integrated is a turnkey cluster certified by Clustercorp Inc. It features a Rackform iServ R266 head node, and iServ R2020 and R2121-IB compute nodes, with Intel® Xeon® Processor 5400 Series technology. Sized to meet workgroup-level needs, scalable to meet department- and enterprise-level needs, featuring leading cluster software, and with a starting configuration price below \$30,000, this is the current cluster configuration of choice.

When you partner with Silicon Mechanics, you get more than high-density, custom-fit cluster solutions—you get an expert like Ken.



For more information about the Hyperform ROCKS+ Integrated cluster visit www.siliconmechanics.com/hpci.

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 in the US and other countries.



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



begins the automatic portion of the installation (Figures 13, 14 and 15). Once installation is complete, the head

node reboots, and you are greeted with your first login screen (Figure 16).



Figure 13. Rocks Installation, 1



Figure 14. Rocks Installation, 2



Figure 15. Rocks Installation, 3

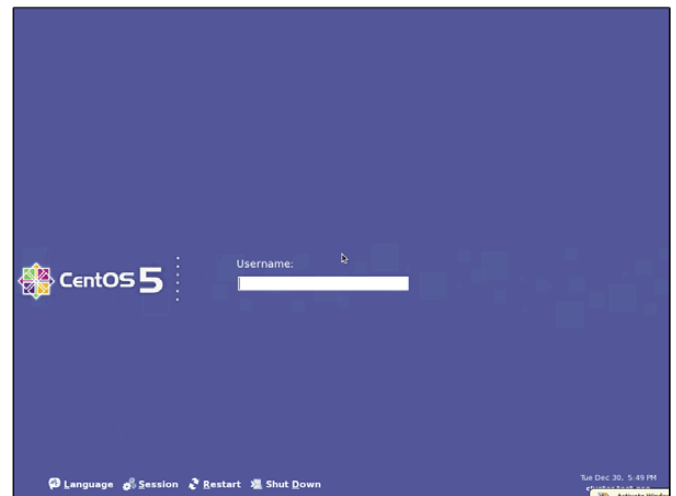


Figure 16. Login Screen

STEP 12 Login

Log on as root, and wait for two or three minutes. This lets the remaining configuration routines finish setting up the cluster in the background. Start a terminal session (Figure 17) to begin installing the compute nodes.

```
login as: root
root@192.168.0.240's password:
Last login: Wed Dec 31 22:17:06 2008 from 192.168.0.181
Rocks 5.1 (V.1)
Profile built 20:34 31-Dec-2008

Kickstarted 12:57 31-Dec-2008
[root@cluster ~]#
```

Figure 17. Root Terminal



All of the systems in the cluster must be synchronized as closely as possible to each other.

STEP 13 Install a Compute Node

Now you're ready to add nodes to the cluster. The Rocks command that accomplishes this is `insert-ethers`. It has quite a few options, but for this example, use the main function of inserting nodes into the cluster. After you invoke `insert-ethers`, you are presented with the screen shown in Figure 18.

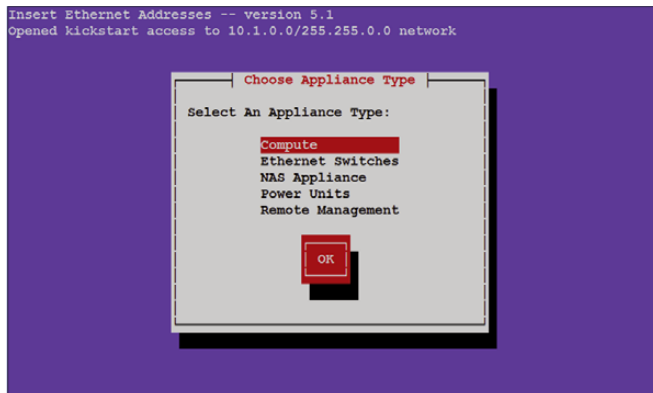


Figure 18. `insert-ethers`

Rocks treats everything that can be connected to the

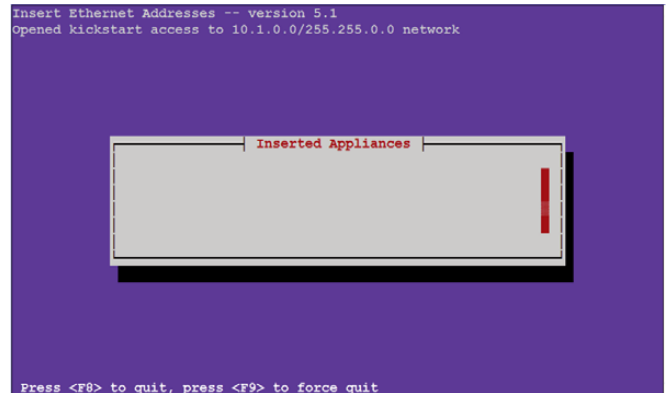


Figure 19. List of Installed Appliances (Empty)

network as an appliance. If it can respond to a command over the network, it's an appliance. For this simple example with a dumb switch, the only things you need to worry about are the compute nodes themselves. Because Compute is already selected, tab to the OK button and press Enter. This brings up the empty list that will be filled with the names and MAC addresses of the nodes as they are added (Figure 19).

Linux News and Headlines Delivered To You

Linux Journal topical RSS feeds NOW AVAILABLE



http://www.linuxjournal.com/rss_feeds

STEP 14 Boot First Compute Node

Now it's time to boot the first compute node. If you have wiped the disk, most systems will start a PXE boot from the network as a default action. If you have a KVM switch and can watch the console on the compute node, you should see the PXE boot begin. When the compute node asks for an address for eth0, you will see the MAC address




Figure 20. List of Inserted Appliances (First Node Added)

STEP 15 Install Additional Compute Nodes

Install the rest of your compute nodes. After a couple more nodes are booted, the list of installed nodes looks like Figure 22.

When the last node in the cluster reboots at the end of its loading process, press F8 on the head node to finish the installation.

Figure 22. List of Installed Appliances (Three Nodes Added)



STEP 16 Quick Functionality Check

Your cluster now is complete and ready for work. First job: roll call. The Rocks cluster-fork function allows the user to execute the same application on all or a subset of the nodes in the cluster. Figure 23 shows executing the uname command via cluster-fork.

The first invocation requires the system to set up the security for each node. Once this is done, subsequent invocations simply run the application. It appears that all of the nodes in the cluster are healthy and ready for work.

If you are looking for a more comprehensive test, take a look at the Intel Cluster Checker package. This application is useful both on a newly created cluster and as a tool for ongoing maintenance.




Figure 21. First Node Installing

entered in the Inserted Appliances list on the head node (Figure 20).

The insert-ethers routine displays the MAC address it has received and the node name it has assigned that node. The () will be filled in with an asterisk (*) when the compute node begins downloading its image (Figure 21).

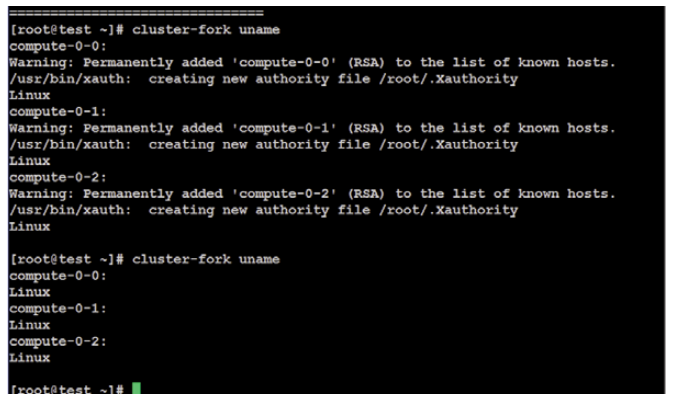



Figure 23. Running uname via cluster-fork

As you enter data on these screens, the installation routine is building a small MySQL database that details all of the component configurations in your cluster.



STEP 17 Install Some Cluster Applications

Now that your cluster is functional, it's time to show it off. One of the more interesting parallel applications is NAMD, a molecular dynamics simulator from the University of Illinois. Paired with VMD, its graphical interface, you essentially have a chemistry set in your cluster.

Summary

When a workstation isn't fast enough, a properly configured cluster can provide all the computing capability you require. Although it is possible to set up a compute cluster manually, many packages are available, both free and commercially supported, that can make the installation and configuration process essentially painless. ■

Tom Lehmann had a 30+ year career with the Intel Corporation working in just about every facet of the computing industry. The last ten years of his Intel career were spent in high-performance computing, in particular, high-performance compute clustering. He is currently a consultant working on system designs, Linux-Windows cluster integration and a Linux-based mainframe emulator called Hercules. Tom, his wife and two dachshunds live in Las Vegas.

Resources

UCSD Rocks: www.rocksclusters.org

Darik's Boot and Nuke: www.dban.org

Intel's Cluster Checker: software.intel.com/en-us/articles/intel-cluster-checker

University of Illinois NAMD: www.ks.uiuc.edu/Research/namd

University of Illinois VMD: www.ks.uiuc.edu/Research/vmd

Advertiser Index

CHECK OUT OUR NEW BUYER'S GUIDE ON-LINE. Go to www.linuxjournal.com/buyersguide where you can learn more about our advertisers or link directly to their Web sites.

Thank you as always for supporting our advertisers by buying their products!

Advertiser	Page #	Advertiser	Page #
1&1 INTERNET, INC. www.oneandone.com	8, 9	POLYWELL COMPUTERS, INC. www.polywell.com	5
ABERDEEN, LLC www.aberdeeninc.com	37	RACKSPACE MANAGED HOSTING www.rackspace.com	C3
ASA COMPUTERS, INC. www.asacomputers.com	23	SAINT ARNOLD BREWING COMPANY www.saintarnold.com	79
CARL.NET www.carl.net	41	SERVERBEACH www.serverbeach.com	13
CORAD, INC. www.coraid.com	25, 79	SERVERS DIRECT www.serversdirect.com	1
DIGI-KEY CORPORATION www.digi-key.com	79	SHAREPOINT www.sptechcon.com	3
EMAC, INC. www.emacinc.com	33, 79	SILICON MECHANICS www.siliconmechanics.com	59, 61
EMPERORLINUX www.emperorlinux.com	21	STRAYTATS www.straytats.com	79
GECAD TECHNOLOGIES/AXIGEN www.axigen.com	79	SUN JAVA ONE CONF. java.sun.com/javaone/index.jsp	35
GENSTOR SYSTEMS, INC. www.genstor.com	67	TECHNOLOGIC SYSTEMS www.embeddedx86.com	77
INTEROP www.interop.com	73	TECHWEB LIVE EVENTS (BLACKHAT) www.blackhat.com	31
LINUX CERTIFIED www.linuxcertified.com	71	UBIQUITI NETWORKS, INC. www.ubnt.com	C2
LOGIC SUPPLY, INC. www.logicsupply.com	19	UTILIKILTS www.utilikilts.com	79
MICROWAY, INC. www.microway.com	C4, 7	ZT GROUP INTERNATIONAL www.ztgroup.com	11
O'REILLY OSCON en.oreilly.com/oscon2009	55		

ATTENTION ADVERTISERS

September 2009 Issue #185 Deadlines
Space Close: June 29; Material Close: July 7

Theme: Cross-Platform Development

BONUS DISTRIBUTIONS:
Utah Open Source Convention, Software Business

Call Joseph Krack to reserve your space
+1-713-344-1956 ext. 118, e-mail joseph@linuxjournal.com

Dojo's Industrial-Strength Grid Widget

Dojo's grid widget provides a fierce alternative to the paginated table that has become such a mainstay of modern-day Web application interfaces. MATTHEW RUSSELL

In almost any application that's complex enough to require a database for persistent storage, there is a need to render data in a tabular format for display or interaction. Although a seemingly simple task, it quickly becomes quite involved once you introduce the need to sort columns, reorder columns, account for editing the contents of cells and managing pagination. As you might expect, this is one of those problems that has been solved so many times that commoditized solutions, such as Dojo's grid widget, now are available for use. After all, you'd much rather concentrate on actually building out the interesting parts of your app than creating and maintaining the infrastructure that holds it all together, right?

If you've been following along with my previous *LJ* articles on Dojo (see Resources), you should have a good idea of how to get Dojo up and running with minimal hassle using AOL's Content Delivery Network (or if you prefer, get it from the official Dojo Web site). As a reminder, the minimal page template is basically an HTML page with a script tag and a `dojo.addOnLoad` block that waits until all cross-domain loading has completed before it executes in order to guarantee that dependencies inside the code block have been met before it executes. You may recall that the skeleton for this page looks something like this:

```
<html>
<head>
  <title>Minimal Development Template</title>
  <script
    type="text/javascript"
    src="http://o.aolcdn.com/dojo/1.2/dojo/dojo.xd.js">
  </script>
  <script type="text/javascript">
    dojo.addOnLoad(function() {
      /* Add Dojo-dependent logic here to
       avoid race conditions */
    });
  </script>
</head>
<body>
</body>
</html>
```

Now I'm going to move to putting the grid widget to work, so consult a reference such as *Dojo: The Definitive Guide* (O'Reilly, June 2008) or on-line documentation at the Dojo Campus if you need a quick refresher.

Note:

Prior to the Dojo 1.2 release, the grid was under heavy development and largely in a state of flux. Although its API still is subject to change, the 1.2 release substantially firmed it up, and aside from accessibility considerations, it is expected to remain intact. Bottom line: don't let the dojo namespace or previous experiences with the grid prior to 1.2 scare you off; it's ready for prime time. Although at the time of this writing, Dojo version 1.3 is just about to be released, the code examples reference version 1.2 and work just fine. There should be minimal, if any, changes necessary to update the examples to version 1.3.

Dojo Data API Primer

The grid widget is necessarily data-centric and builds directly upon the abstractions offered by the toolkit's robust data APIs, so a very brief introduction is helpful for setting that context. In Dojo parlance, the abstraction that the data APIs offer is a store that contains items, where a store can implement a particular subset of the four `dojo.data` APIs:

- Read: APIs for querying items in a store.
- Identify: APIs for uniquely identifying items in a store.
- Write: APIs for creating, modifying and deleting items in a store.
- Notification: APIs for triggering event handlers when items in a store are created, modified or deleted.

The toolkit comes stocked with two handy store implementations that are bundled into the `dojo.data` module: the `ItemFileReadStore` and the `ItemFileWriteStore`. The `ItemFileReadStore` implements the Read and Identity APIs, while the `ItemFileWriteStore` implements all four of these APIs. Be aware, however, that the `dojox.data` module contains a plethora of additional data modules for common tasks that you *will* want to leverage to your advantage; interfacing with comma-separated value (CSV) files, Flickr, Amazon's S3 service, OPML files and Atom content are



Linux - FreeBSD - x86 Solaris - MS etc.



Proven technology. Proven reliability.

When you can't afford to take chances with your business data or productivity, rely on a GS-1245 Server powered by the Intel® Xeon® Processors.

Quad Core Woodcrest



2 Nodes & Up to 16 Cores - in 1U

Ideal for high density clustering in standard 1U form factor. Upto 16 Cores for high CPU needs. Easy to configure failover nodes.

Features:

- 1U rack-optimized chassis (1.75in.)
- Up to 2 Quad Core Intel® Xeon® Woodcrest per Node with 1600 MHz system bus
- Up to 16 Woodcrest Cores Per 1U rackspace
- Up to 64GB DDR2.667 & 533 SDRAM Fully Buffered DIMM (FB-DIMM) Per Node
- Dual-port Gigabit Ethernet Per Node
- 2 SATA Removable HDD Per Node
- 1 (x8) PCI_Express Per Node



Servers :: Storage :: Appliances

Genstor Systems, Inc.

780 Montague Express. # 604
San Jose, CA 95131

Www.genstor.com

Email: sales@genstor.com

Phone: 1-877-25 SERVER or 1-408-383-0120



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

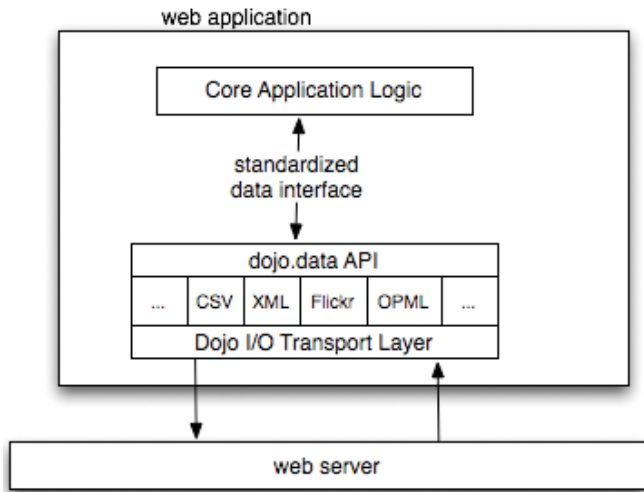


Figure 1. Building on top of an abstraction such as the dojo.data layer thins out the application logic and provides a uniform abstraction for accessing data that is more maintainable and much less brittle.

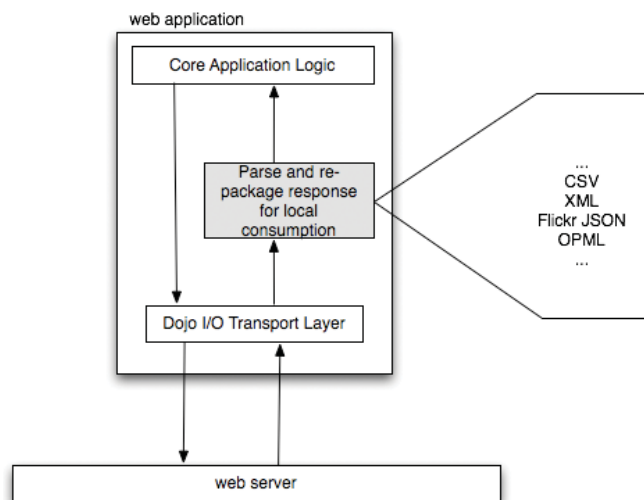


Figure 2. Retrieving remote data from a server and massaging it into the format expected by client-side JavaScript code has been a bane of Web development for some time, and generally produces brittle designs that do not sufficiently decouple the application logic from the underlying data format being used.

just a few of the handy implementations available (Figures 1 and 2).

In general, you provide data to an ItemFileReadStore or ItemFileWriteStore in one of three ways:

- By feeding it a file containing JSON data that meets a particular structure.
- By feeding it a JavaScript object conforming to the same structure as the JSON file data.

- By programmatically creating new items for the store.

Let's briefly consider examples that illustrate these operations. The grid widget simply reflects whatever data is in the store that backs it, so understanding how to manipulate data stores is essential to controlling what ultimately appears in a grid widget (Listing 1).

Although the example in Listing 1 illustrates providing inline JavaScript data for the store to consume, the same data could just as easily have been provided by way of an I/O request to the server. For example, the store could have taken a URL parameter on construction, which would have fetched a file. Assuming its contents were the same JavaScript object identified by the data property in the previous example, the results would have been the same:

```
/* Another way to create a store */
var store = new dojo.data.ItemFileReadStore({
  url : "/some/server/side/url"
});
```

For more information on the dojo.data API, consult Chapter 9 of *Dojo: The Definitive Guide*, on-line documentation available at the Dojo Campus, or read the API well-documented

specs that are packaged in the Dojo source code itself in the dojo.data.api namespace.

Listing 1. ItemFileWriteStore Example

```
<html>
<head>
  <title>ItemFileWriteStore Example</title>
  <script
    type="text/javascript"
    src="http://o.aolcdn.com/dojo/1.2/dojo/dojo.xd.js">
  </script>
  <script type="text/javascript">
    dojo.require("dojo.data.ItemFileWriteStore");
    dojo.addOnLoad(
      function() {

        /* Creating a store with inline JavaScript data */
        var store = new dojo.data.ItemFileWriteStore({
          data : {
            identifier : "id",
            items : [
              {"id" : 1, "label" : "foo"},
              {"id" : 2, "label" : "bar"},
              {"id" : 3, "label" : "baz"}
            ]
          }
        });

        /* Add another item - a synchronous operation */
        store newItem({"id" : 4, "label" : "qux"});

        /* Fetch an item with id=4 - an asynchronous operation */
        store.fetch({
          query : {id : 4},
          onItem : function(item) {
            console.log("Asynchronous callback for fetching item:",
              item);
          }
        });

        /* Delete the item with id=4, a synchronous operation */
        store.deleteItem(item);

        /* Save the results */
        store.save();

        /* Could have reverted the results with
           store.revert(); */
      }
    });
  </script>
</head>
<body>
</body>
</html>
```

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

Grid Fundamentals

With newly found knowledge of how to create and manipulate data stores, we're now ready to bind one of those stores to the grid widget. Listing 2 is a full-blown example of programmatically creating a simple DataGrid and attaching an ItemFileReadStore.

Additions to Listing 2 include adding some CSS files to style the grid and a few extra lines of script to specify column information. The final call to startup() is a fairly standard Dijit life-cycle method that needs to be called to tell widgets to lay themselves out when they are constructed programmatically.

Now, let's turn to creating a grid widget in markup. As you're about to see, creating a grid in markup is as simple as defining an HTML table structure (Listing 3).

Aside from declaring the ItemFileReadStore and DataGrid in markup, the only other change to note is that the parseOnLoad configuration switch was provided to the SCRIPT tag that loads Dojo.

The dojo.parser module also was included as a dependency, because it's what actually scans the BODY of the page for

ID	Label
1	foo
2	bar
3	baz

Figure 3. A simple grid created from Listing 2: column sorting is built in by default.

dojoType tags and instantiates any widgets that are found. Depending on your programming background and your project's overall design, you may prefer markup-driven development to a script-driven approach. Dojo provides facilities for you to accomplish the very same objectives either way, so you're covered in either case and have the flexibility to choose.

Listing 2. Simple Data Grid

```
<html>
<head>
  <title>Simple Data Grid</title>

  <link rel="stylesheet" type="text/css" href=
    "http://o.aolcdn.com/dojo/1.2/dojo/resources/dojo.css"/>
  <link rel="stylesheet" type="text/css" href=
    "http://o.aolcdn.com/dojo/1.2/dijit/themes/tundra/tundra.css"/>
  <link rel="stylesheet" type="text/css" href=
    "http://o.aolcdn.com/dojo/1.2/dojox/grid/resources/Grid.css"/>
  <link rel="stylesheet" type="text/css" href=
    "http://o.aolcdn.com/dojo/1.2/dojox/grid/resources/tundraGrid.css"/>

  <style type="text/css">
    #gridNode {
      width: 200px;
      height: 200px;
    }
  </style>

  <script type="text/javascript"
    src="http://o.aolcdn.com/dojo/1.2/dojo/dojo.xd.js">
  </script>

  <script type="text/javascript">
    dojo.require("dojo.data.ItemFileReadStore");
    dojo.require("dojox.grid.DataGrid");

    dojo.addOnLoad(
      function() {
        /* Create a store with inline JavaScript data */
        var gridStore = new dojo.data.ItemFileReadStore({
          data : {
            identifier : "id",
            items : [
              {"id" : 1, "label" : "foo"},
              {"id" : 2, "label" : "bar"},
              {"id" : 3, "label" : "baz"}
            ]
          }
        });

        /* A simple layout that specifies column headers and
         * mappings to fields in the store */
        var gridLayout = [
          {name : "ID", field : "id", width : "50%"},
          {name : "Label", field : "label", width : "50%"}
        ];

        /* Programmatically construct a data grid */
        var grid = new dojox.grid.DataGrid({
          store : gridStore,
          structure : gridLayout
        }, "gridNode");

        /* Tell the grid to lay itself out since
         * it was programmatically constructed */
        grid.startup();
      }
    );
  </script>
</head>
<body class="tundra">
  <div id="gridNode"></div>
</body>
</html>
```

Managing the Selection

A discussion of every possible DataGrid property or method is well beyond the scope of this article, but it's worthwhile to walk through a few of the more common operations you'll likely need to get up and running. Once you have a grid loaded with data, it's likely that one of the first things you'll want to do is determine what is selected and access the data contained in the selection. The DataGrid exposes a property called selection that is a fairly sophisticated Object providing the key methods for retrieving and manipulating the selection. Recalling that the way to retrieve a reference to a widget is through the dijit.byId method, you can gain access to the selection from Listing 3 simply by calling dijit.byId("gridNode").selection. Using Firebug or consulting the source code or on-line documentation, you would discover a number of useful properties. A few of the most commonly used include:

- `getSelected()`: returns an array of `dojo.data` items that are reflected by the current selection.
- `select(/*Integer*/ idx)`: sets the current selection to the row index identified.
- `deselect(/*Integer*/ idx)`: removes the row index from the current selection.

- `selectRange(/*Integer*/startIdx, /*Integer*/ endIdx)`: selects the rows identified by the start and end indexes, inclusively.
- `clear()`: clears the selection.
- `onSelected(/*Integer*/ idx)`: an extension point that can be overridden to supply custom functionality whenever a particular row is selected.
- `onDeselected(/*Integer*/ idx)`: an extension point that can be overridden to supply custom functionality whenever a particular row is deselected.

Consider the following examples to get an idea of how you might put the DataGrid's selection property to work:

```
/* Assume a grid identified by a node with id=gridNode
 * that has lots of rows and no selection */

/* select rows 11-20 inclusive */
dijit.byId("gridNode").selection.selectRange(11,20);

/* Attach a custom event handler for row selection */
dijit.byId("gridNode").selection.onSelected = function(idx) {
    console.log("onSelected", idx);
};
```

Listing 3. Simple Data Grid

```
<html>
<head>
  <title>Simple Data Grid</title>

  <link rel="stylesheet" type="text/css" href=
    "http://o.aolcdn.com/dojo/1.2/dojo/resources/dojo.css"/>
  <link rel="stylesheet" type="text/css" href=
    "http://o.aolcdn.com/dojo/1.2/dijit/themes/tundra/tundra.css"/>
  <link rel="stylesheet" type="text/css" href=
    "http://o.aolcdn.com/dojo/1.2/dojox/grid/resources/Grid.css"/>
  <link rel="stylesheet" type="text/css" href=
    "http://o.aolcdn.com/dojo/1.2/dojox/grid/resources/tundraGrid.css"/>

  <style type="text/css">
    #gridNode {
      width: 200px;
      height: 200px;
    }
  </style>

  <script type="text/javascript"
    src="http://o.aolcdn.com/dojo/1.2/dojo/dojo.xd.js"
    djConfig="parseOnLoad:true">
  </script>

  <script type="text/javascript">
    dojo.require("dojo.data.ItemFileReadStore");
    dojo.require("dojox.grid.DataGrid");
    dojo.require("dojo.parser");
  </script>
</head>
<body class="tundra">
  <!-- Fetch data from a store -->
  <span dojoType="dojo.data.ItemFileReadStore"
    jsId="gridStore"
    url="data.json">
  </span>

  <!-- Define the grid directly in markup and allow the parser
    to take care of the rest -->
  <table id="gridNode" dojoType="dojox.grid.DataGrid" store="gridStore">
    <thead>
      <tr>
        <th width="50%" field="id">ID</th>
        <th width="50%" field="label">Label</th>
      </tr>
    </thead>
  </table>
</body>
</html>
```



```

/* Clears the selection and reselects only row 13.
 * Note that the onSelect event handler fires. */
dijit.byId("gridNode").selection.select(13);

```

As with many OSS projects, the only authoritative API reference is the source code itself or the documentation generated directly from it, so check there for a complete listing and for more details.

Common Grid Operations

Although the DataGrid's selection property provides useful methods, such as `getSelection`, that facilitate accessing `dojo.data` items that back the interface, the DataGrid itself offers myriad methods of its own that provide more direct and finely grained access to the grid and the data that backs it. Here's a quick synopsis of just a few common ones:

- `getItem(/*Integer*/ idx)`: returns a `dojo.data` item reflected in the row index.
- `setQuery(/*Object*/ query, /*Object*/ queryOptions)`: filters the data in the table by executing the query and query options against the store that backs the grid.
- `setStore(/*Object*/ store, /*Object*/ query, /*Object*/ queryOptions)`: disposes of references to the existing store and attaches a new one, optionally passing in a query and query options for filtering.
- `onRowClick(/*Event*/ evt)`: called when a cell is clicked; `evt` is a decorated W3C Event object.
- `onCellClick(/*Event*/ evt)`: called when a cell is clicked; `evt` is a decorated W3C Event object.
- `onCellFocus(/*Object*/ cell, /*Integer*/ rowIdx)`: called when a cell receives focus.
- `setStructure(/*Object|Array*/ structure)`: provides a row of changing the grid's layout after it is initially rendered.
- `scrollToRow(/*Index*/ idx)`: scrolls the grid to the row index.
- `setSortInfo(/*Object*/ obj)`: called to set sorting criteria.
- `sort()`: sorts the grid according to the information supplied by `setSortInfo`.
- `columnReordering`: a property that allows for drag-and-drop column reordering on the grid.

As usual, once the API has been unearthed, the implementation details of putting the grid to work are usually straightforward enough. Here are a few examples to get the wheels turning:

```

/* Filter the grid such that only row items having
 * a name that starts with the letter B appear */
dijit.byId("gridNode").setQuery({name : "B*"});

/* Get the item reflected in row 23 */
dijit.byId("gridNode").getItem(23);

dijit.byId("gridNode").onRowClick = function(evt) {
    /* Display interesting parts of the decorated Event Object */
    console.log("onRowClick: cell", evt.cell);
    console.log("onRowClick: cellIndex", evt.cellIndex);
    console.log("onRowClick: row", evt.row);
    console.log("onRowClick: rowIndex", evt.rowIndex);
    console.log("onRowClick: grid", evt.grid);
};

```

Editable Data

Given that a grid widget often is nothing more than a visual interface into a data store, it won't be long before you'll not only want to view the data in the store, but also edit it and persist it back to the server. A great testament to the flexibility



Linux Certified

Linux Laptops: The New LC2000 Series



- High Performance
- Amazing ROI
- Robust
- Fully Compatible
- Cost Effective

Open Source Training, Services and Products 1-877-800-6873 www.linuxcertified.com

of Dojo's data APIs is how the grid's architecture builds directly upon the Write and Notification APIs in particular to make this as easy as it should be—meaning, so long as your store implements Read, Identity, Write and Notification, a grid that you attach to it is capable of providing editable cells that “just work”. In other words, the reason you can attach a stock component like an ItemFileWriteStore to the DataGrid and get editable data isn't because of specialized logic that binds the two together, but simply because the ItemFileWriteStore

implements the full spectrum of dojo.data APIs. Although the ItemFileReadStore is used for pedagogical purposes, numerous highly useful store implementations are included as part of the dojo.data module.

The necessary changes to the minimal example we have been working with in the previous Listings are surprisingly simple; just change the store to an implementation that supports the Write and Notification APIs and provide a couple extra attributes in the markup for any columns that should be editable. The

Listing 4. An ultra-simple Web server that provides slices of a very large (mock) data source for a dojo.grid.Grid client that uses a dojo.data.QueryReadStore to page the data on demand.

```
import cherrypy #do an "easy_install cherrypy" to get it
from cherrypy.lib.static import serve_file

import demjson #do an "easy_install demjson" to get it
import os
from random import randint #for building up mock data

json = demjson.JSON(compactly=False)
jsonify = json.encode

NUM_ITEMS = 1000000

class Content:
    def __init__(self):
        """
        Maybe you would call out to a db with some sql to get some
        data based on the query string that comes into /data. For
        now, we'll build up some static data to use.
        """

        self.items = []

        possible_item_labels = ["foo", "bar", "baz", "qux"]
        id=0
        for i in xrange(NUM_ITEMS):
            self.items.append({
                'id' : id,
                'label' : possible_item_labels[randint(0,3)]
            })
            id +=1

        #keep track of sort order b/c sorting is expensive...
        self.current_sort_order = ""

    @cherrypy.expose
    def data(self, **kw):
        """
        Serve up the data via http://localhost:8000/data

        kw will contain whatever is in your store's query.
        By default the query string will come across as
        something like:
        ?name=*&start=0&count=20 to populate the table

        Note: you may get into trouble if you have multiple users
        trying to access this url and changing the sort order of
        items all at the same time (but relax, this is just
        a little demo.)
        """

        #sorting the items by values for a given dictionary key...
        if kw.get('sort') and self.current_sort_order != kw.get('sort'):
            if kw['sort'][0] == '-': #descending order, slice off the -
                self.items.sort(lambda m,n:cmp(m.get(kw['sort'])[1:]), \
                    n.get(kw['sort'])[1:]),reverse=True)
            else: #ascending order
                self.items.sort(lambda m,n:cmp(m.get(kw['sort']), \
                    n.get(kw['sort'])))
            self.current_sort_order = kw['sort']

        #slicing the data...
        start = int(kw['start'])
        end = start + int(kw['count'])

        #serving up the slice of interest as well as the total size
        return jsonify({
            'numRows':NUM_ITEMS,
            'items':self.items[start:end],
            'identifier' : 'id'
        })

    @cherrypy.expose
    def index(self, **kw):
        """
        Serve up the web page through http://localhost:8000/
        """
        return serve_file(os.path.join(os.getcwd(), 'page.html'))
        #the page containing the grid

cherrypy.server.socket_port = 8000
cherrypy.quickstart(Content(), '/')
```

How can **mobile solutions**
help us increase revenue?

How will **cloud computing** help
us do more with less?

What technologies will help us cut costs...quickly?

How can **virtualization** make
our business more efficient?

Will **unified communications**
help us improve customer
responsiveness?



INTEROP[®]

LAS VEGAS | MAY 17–21
Mandalay Bay Convention Center

REGISTER TODAY
to save up to \$500 on
Conference Passes or reserve
your free Expo Pass.
www.interop.com/lv

ANSWERS. ACTION. ADVANTAGE.

Get the information you need to build your competitive edge at IT's most comprehensive Conference and Expo. See all the latest innovations—including virtualization, mobility and cloud computing—that will help you increase efficiency, drive revenue and cut costs.

Attend 200+ sessions and see 400+ exhibitors at the leading business technology event.

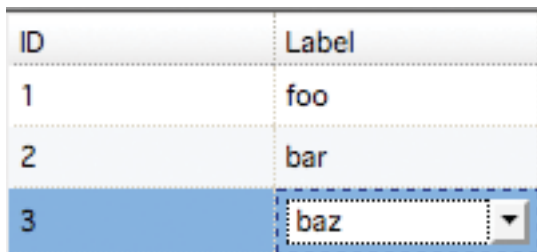
Register with priority code CMXZNL05 by April 17 to save up to \$500 on conference passes or reserve your free expo pass. www.interop.com/lv

dojox.grid.cells module provides lightweight wrappers around many common form widgets from Dijit, so let's take a look at a simple change that would make the Label column editable by introducing a select box:

```
<body class="tundra">
  <!--Remember to have dojo.require'd the ItemFileWriteStore -->
  <span dojoType="dojo.data.ItemFileWriteStore"
    jsId="gridStore"
    url="data.json">
  </span>

  <table id="gridNode"
    dojoType="dojox.grid.DataGrid"
    store="gridStore">
    <thead>
      <tr>
        <th width="50%" field="id">ID</th>
        <th width="50%" field="label"
          cellType="dojox.grid.cells.Select"
          options="foo,bar,baz,qux"
          editable="true">Label</th>
      </tr>
    </thead>
  </table>
</body>
```

As you might imagine, the DataGrid's custom event handlers, such as onFocus, onBlur, onApplyEdit, onCancelEdit and so on, become increasingly useful for an editable interface. As always, you also may attach event handlers to the store that backs the grid if handling changes at the data level seems more appropriate for your application than tracking UI-related events. Figure 4 shows the editable grid.



ID	Label
1	foo
2	bar
3	baz

Figure 4. An Example of an Editable Interface in the DataGrid

Server-Backed Data

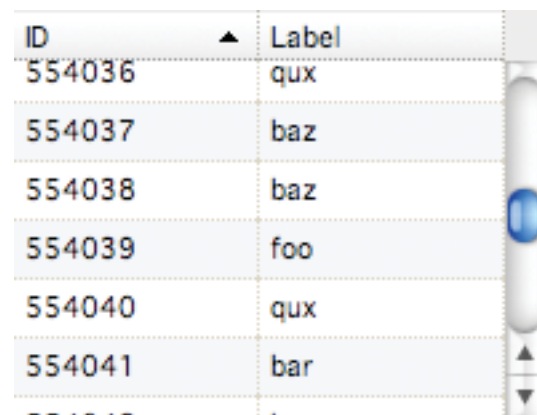
So far, the examples demonstrated here are using ItemFileReadStore or ItemFileWriteStore, which necessarily implies that your data set is small enough that it's practical to load it into the client. In other words, we've been dodging the issue of having such a large data set (say, millions of records) that it can't all be loaded into the client. Let's put together a final example that demonstrates the grid at work using a server-backed store, such as the dojox.data.QueryReadStore. The markup for defining the DataGrid should look familiar enough. Note that because the QueryReadStore implements only the

Read and Identity APIs, trying to make cells editable would have no effect. It's totally possible, however, to extend the QueryReadStore with Write and Notification support or attach a store, such as the dojox.data.JsonRestStore, that implements all four dojox.data APIs to produce an editable interface:

```
<body class="tundra">
  <!--Fetch data from a store as usual.
  This time, it just happens to be a QueryReadStore -->
  <span dojoType="dojox.data.QueryReadStore"
    jsId="gridStore"
    url="/data">
  </span>

  <!-- Define the grid directly in markup and allow the parser
  to take care of the rest -->
  <table id="gridNode"
    dojoType="dojox.grid.DataGrid" store="gridStore">
    <thead>
      <tr>
        <th width="50%" field="id">ID</th>
        <th width="50%" field="label">Label</th>
      </tr>
    </thead>
  </table>
</body>
```

To try out the example, however, you need a basic server implementation that returns pages of data whenever the QueryReadStore requests them. A minimalistic server written in CherryPy is shown in Listing 4.



ID	Label
554036	qux
554037	baz
554038	baz
554039	foo
554040	qux
554041	bar

Figure 5. Given a server-backed store, the DataGrid can render arbitrary numbers of rows—all without pagination!

Wrapping Up

Although we barely scratched the surface of the DataGrid widget's utility or power, you hopefully have a good feel for some of the things you can do with it. Dojo's source code includes a number of useful examples that are bundled as tests, and they provide a great way to get rolling with more grid goodness. You also might drop by the

#dojo IRC room on freenode.net to get some help or share what you're doing. ■

Matthew Russell is an open Web technology consultant and the author of *Dojo: The Definitive Guide* (O'Reilly, June 2008).

Resources

"Dojo: the JavaScript Toolkit with Industrial-Strength Mojo" by Matthew Russell, *LJ*, July 2008: www.linuxjournal.com/article/9900

"Dojo, Now with Drawing Tools!" by Matthew Russell, *LJ*, February 2009: www.linuxjournal.com/article/10308

The Dojo Toolkit: dojotoolkit.org

Dojo Campus: dojocampus.org

Dojo Campus Data Tutorial: docs.dojocampus.org/quickstart/data/usingdatastores

Official Dojo API: api.dojotoolkit.org

Official dojox.grid.DataGrid API: api.dojotoolkit.org/jsdoc/dojox/1.2/dojox.grid

Lots of Working Grid Examples: archive.dojotoolkit.org/nightly/dojotoolkit/dojox/grid/tests

Important Grid Bug for Firefox Users: bugs.dojotoolkit.org/ticket/8242

ON THE WEB, ARTICLES TALK!

Phil Zimmermann comes to the rescue again with ZRTP, a protocol for securely transferring keys across a VOIP network: www.linuxjournal.com/video/secure-voip-communication-zrtp.



SNMP Monitoring with Nagios

Using Nagios, you can monitor Dell servers with SNMP via Dell's server administration tools. JASON ELLISON

Nagios has been around since 2002 and is considered stable software. It is in use by the likes of American Public Media, JP Morgan Chase and Yahoo, just to name a few. It is an enterprise-level network and systems-monitoring platform. Nagios performs checks of services and hosts using external programs called Nagios plugins.

SNMP (Simple Network Management Protocol) is a network protocol designed for monitoring network-attached devices. It uses OIDs (Object Identifiers) for defining the information, known as MIBs (Management Information Base), that can be monitored. The design is extensible, so vendors can define their own items to be monitored.

OpenManage is provided with Dell servers and is an extremely well-documented system (see Resources) that provides extensive server administration capabilities. OpenManage works with both Linux and Windows. The OpenManage "SNMP Reference Guide" (see Resources) is a 732-page document that is "intended for system administrators, network administrators and anyone who wants to write SNMP MIB applications to monitor systems". The "SNMP Reference Guide" documents the SNMP OIDs/MIBs for monitoring Dell's servers.

The system described here was implemented for a local utility company when it upgraded to Dell Power Edge servers. As often is the case, out of the box, Nagios didn't do exactly what the company needed, but being an open-source project, it easily was extended to accomplish the goal. All we needed was a Nagios plugin to monitor the new servers.

Don't Re-invent the Wheel

The first thing I set out to do was find an existing Nagios plugin that offered similar functionality to what we needed. Quite a number of existing plugins are available. In

The system described here was implemented for a local utility company when it upgraded to Dell Power Edge servers.

less than one hour, I found `check_snmp_temperature.pl` by William Leibzon. This is a plugin module that monitors the temperature of various devices remotely via SNMP. Although monitoring temperatures was not our goal, retrieving information via SNMP and reporting it to Nagios was. The module is written in Perl and after reading it over, it looked very well written.

Chapter 4 of the Dell's "SNMP Reference Guide" is the "System State Group". It states:

The Management Information Base (MIB) variables presented in this section enable you to track various attributes that describe the state of the critical components supported by your system. Components monitored under the System State Group include power supplies, AC power cords, AC power switches, and cooling devices, as well as temperature, fan, amperage, and voltage probes.

The associated OIDs provide the overall state of all the critical subsystems that we were interested in. OIDs exist that provide much greater detail, but in this situation, the requirement was to be alerted only if a server had a problem and to indicate the particular subsystem that had the problem.

One of the benefits to choosing these particular OIDs turned out to be that they all respond in the same format.

One subsystem was not addressed in the "System State Group" chapter—the RAID subsystem. There is, however, an OID for monitoring it. This OID is described in Chapter 23, the "Storage Management Group".

As stated earlier, these OIDs are used to define particular MIBs that can be queried via SNMP. On the Dell server, there is an SNMP server running. The SNMP server answers queries that are in the form of a long string of numbers (the OID). This string of numbers is understood by the SNMP server to be a specific question. For instance, if you want to ask the SNMP server "How are your power supplies?", you would send it the OID `.1.3.6.1.4.1.674.10892.1.200.10.1.9.1` (Figure 1). The SNMP server will respond with 3 if the power supplies are okay.

Table 1 shows the OIDs we are interested in.

One of the benefits to choosing these particular OIDs turned out to be that they all respond in the same format. Dell refers to this format as `DellStatus`, and it maps integers to subsystem states:

Variable Name:	<code>DellStatus</code>
Data Type:	Integer
Possible Data Values	Meaning of Data Value:
<code>other(1)</code>	The object's status is not one of the following:

<code>unknown(2)</code>	The object's status is unknown.
<code>ok(3)</code>	The object's status is OK.
<code>nonCritical(4)</code>	The object's status is warning, noncritical.
<code>critical(5)</code>	The object's status is critical (failure).
<code>nonRecoverable(6)</code>	The object's status is nonrecoverable (dead).

Now that we knew what we wanted to monitor, it was time to modify `check_snmp_temperature.pl` to do what was needed. The result, `check_dell_openmanager.0.7-test.pl`, is too long to print here, but it is available on the *Linux Journal* FTP site (see Resources).

Testing

Because I did not have a spare Dell Power Edge server sitting around to test the modified script, I had to test it another way. Reading the man page for `snmpd.conf`, I found that you could have external programs answer certain OIDs using "pass-through" scripts. The bash script (`dell_open_manager_test.sh`) below serves as my pass-through script for testing. With this script, I can simulate all of the states that the Dell server could be in:

```
#!/bin/bash
#
# bash script to replicate a working Dell OpenManage SNMP agent
# works with Net-SNMP daemon. infotek@gmail.com
#

REQUEST_OID="$2"

echo "$REQUEST_OID";
case "$REQUEST_OID" in
    .1.3.6.1.4.1.674.10892.1.200.10.1.4.1)
        echo "integer"; echo "3"; exit 0 ;;
    .1.3.6.1.4.1.674.10892.1.200.10.1.9.1)
        echo "integer"; echo "5"; exit 0 ;;
    .1.3.6.1.4.1.674.10892.1.200.10.1.12.1)
        echo "integer"; echo "3"; exit 0 ;;
    .1.3.6.1.4.1.674.10892.1.200.10.1.21.1)
        echo "integer"; echo "4"; exit 0 ;;
    .1.3.6.1.4.1.674.10892.1.200.10.1.24.1)
        echo "integer"; echo "3"; exit 0 ;;
    .1.3.6.1.4.1.674.10892.1.200.10.1.27.1)
        echo "integer"; echo "3"; exit 0 ;;
    .1.3.6.1.4.1.674.10892.1.200.10.1.30.1)
        echo "integer"; echo "3"; exit 0 ;;
    .1.3.6.1.4.1.674.10892.1.200.10.1.41.1)
        echo "integer"; echo "3"; exit 0 ;;
    .1.3.6.1.4.1.674.10893.1.20.110.13.0)
        echo "integer"; echo "3"; exit 0 ;;
    *)
        echo "string"; echo "$@"; exit 0 ;;
esac
exit
```

To use the script, I added the following lines to the end of `/etc/snmp/snmpd.conf`:


```
### dell open manager test
view systemview included .1.3.6.1.4.1.674
pass .1.3.6.1.4.1.674 /bin/bash \
    /usr/local/bin/dell_open_manager_test.sh
```

To make the changes in the configuration file take effect, restart the snmpd daemon. On Slackware, this is done via the following:

```
# /etc/rc.d/rc.snmpd restart
Shutting down snmpd: . DONE
Starting snmpd: /usr/sbin/snmpd -A -p \
    /var/run/snmpd -a -c /etc/snmp/snmpd.conf
```

To query the SNMP server, we use Net-SNMP's command-line snmpget utility:

```
# snmpget -v 1 -c public 127.0.0.1 \
    .1.3.6.1.4.1.674.10892.1.200.10.1.9.1
SNMPv2-SMI::enterprises.674.10892.1.200.10.1.9.1 = INTEGER: 3
```

The response is an integer value of 3. The value 3 in the DellStatus (see above) maps to "ok(3)". The object's status is "OK". This tells us that the pass-through script is working. Now, we test the /check_dell_openmanager.pl Perl script:

```
# ./check_dell_openmanager.pl -H 127.0.0.1 -C public -T pe2950
OK
```

To test other values, simply modify the dell_open_manager_test.sh shell script. For example, to simulate an error in the Cooling Device OID (.1.3.6.1.4.1.674.10892.1.200.10.1.21), modify that OID's line in the script to return a code of 4 for nonCritical:

```
.1.3.6.1.4.1.674.10892.1.200.10.1.21.1)
    echo "integer"; echo "4"; exit 0 ; ;
```

Now, running the Perl script produces a warning:

```
# ./check_dell_openmanager.pl -H 127.0.0.1 -C public -T pe2950
WARNING:Cooling Device Status=Non-Critical
```

To simulate a critical error, let's modify the Power Supply OID to reply with a 5:

```
.1.3.6.1.4.1.674.10892.1.200.10.1.9.1)
    echo "integer"; echo "5"; exit 0 ; ;
```

```
# ./check_dell_openmanager.pl -H 127.0.0.1 -C public -T pe2950
CRITICAL:Cooling Device Status=Non-Critical, \
    Power Supply Status=Critical
```

To test the script on the live production systems, we added the check_dell_openmanager.pl command to a working Nagios server. We opened the case cover on a live system to generate a Chassis Intrusion Status error to test the plugin.

Within a few seconds, we had an SMS message on the IT administrator's phone letting us know that there was a problem with the chassis subsystem on the server we just opened.

Moving Forward

After writing this plugin, I uploaded it to a Web site that hosts third-party addons for Nagios named Nagios Exchange. In short order, I was getting e-mail messages from all over the world concerning the Nagios plugin I had written. Some were suggestions, and some were from people in need of help. It was not an overwhelming number of messages. At most, two a week and sometimes none. It was just enough to let me know that people other than me actually were using this thing.

I would like to make a few improvements to the module. For one, I think there may be a way to reduce the SNMP queries to only one query to obtain the overall global status of the machine. Then, only if the state is not "ok(3)", move to query the other OIDs so that a more specific error can be reported.

It also would be nice to be able to evaluate the existence of the various subsystems, that way, for example, if a machine has a RAID array, it is monitored, and if not, the script skips it.

One of the most common e-mail messages I get is about missing the Net::SNMP Perl module. I would like to test for these common-case scenarios. If the test fails, I would like to print the problem with a solution. In the case of "Net::SNMP", it should print:

```
You are missing the Net::SNMP perl module.
Please install it using:
perl -MCPAN -e shell
cpan> install "Net::SNMP"
```

This would improve end-user experience significantly, especially for users new to Linux. ■

Jason Ellison (www.jasonellison.net) is an independent IT consultant in the Gulf Coast area. He can be reached at infotek@gmail.com.

Resources

check_dell_openmanager.0.7-test.pl: ftp.linuxjournal.com/pub/lj/listings/issue182/10204.tgz

Nagios: www.nagios.org

check_dell_openmanager: www.nagiosexchange.org/cgi-bin/page.cgi?g=1437.html

Dell OpenManage Server Administrator Documentation: support.dell.com/support/edocs/software/svradmin

Dell OpenManage Server Administrator Version 5.4 SNMP Reference Guide: support.dell.com/support/edocs/software/svradmin/5.4/en/snmp/pdf/om_54_snmp_ref_gd.pdf



GET NOTICED!

With custom, temporary tattoos! Any image!

5% Off With This Code:
linuxjournal

Great for conventions!

 **StrayTats.com**



axigen
MAIL SERVER

The Mail Server
for IT Professionals

www.axigen.com

TEXAS' OLDEST CRAFT BREWERY



WWW.SAINTARNOLD.COM

EtherDrive®

The AFFORDABLE Network Storage

Fiber Channel speeds at Ethernet prices

CORAIID |  **vmware** | technology alliance PARTNER
ESX 3.5 compatible EtherDrive® HBA

INNOVATION ON THE GO
ORDER YOUR BEAGLE BOARD FROM DIGIKEY.COM



AVAILABLE EXCLUSIVELY AT DIGI-KEY
 **beagleboard**

only \$149⁰⁰


**LOW-COST, NO FAN,
SINGLE-BOARD
COMPUTER**

Digi-Key
CORPORATION


www.digikey.com

ARM9 System on Module
Internet Appliance Engine **SoM-9G20**

- Atmel ARM9 400Mhz CPU
- 10/100 BaseT Ethernet
- SD/MMC Flash Card Interface
- 2 USB 2.0 Host Ports & 1 Device Port
- 6 Serial Ports, 2 SPIs & Audio Interface



The SoM-9G20 is the ideal processor engine for your next design. The System on Module (SoM) approach provides the flexibility of a fully customized product at a greatly reduced cost. Single unit pricing starts at \$155.

 **EMAC Linux 2.6 Kernel**

EMAC, inc.
EQUIPMENT MONITOR AND CONTROL

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

American made Utility Kilts for Everyday Wear

UTILIKILTS.com



The Post-Monopoly Game

Walking out of cable TV's walled garden. DOC SEARLS



Far as I know (that is, as far as some Motorola engineers have told me), my Dish Network and Verizon FiOS set-top boxes are Linux machines. So is my Sony flat-screen TV, which came complete with a four-page document explaining the GPL.

Linux in each case is embedded. That is, it is enslaved to a single purpose or to a narrow set of purposes. This isn't a big deal. Linux has become the default embedded operating system for all kinds of stuff. I just think TV would have a lot bigger future if we liberated the whole category from enslavement to Hollywood and its captive distributors.

Until we do, the one-way-ness of TV remains a highway to hell.

I'm getting a good look at that hell right now, sitting in an airport lounge in Boston. It's still winter as I'm writing this. There are lots of canceled flights and, therefore, lots of relevant news on the lounge television, tuned, as always, to CNN. If this were two years ago, there would have been people gathered around that TV to see what's up with the weather and the rest of the news.

But, not today. Nobody is watching. The TV is just noise in the background. Of the 18 passengers waiting here, all but two are using laptops. I just did a quick walk-around and talked to a few of the laptopers. All of them are using their laptops to keep up with weather and flight conditions. TV can't compete with that. There are too many good sources of information on the Web now. More important, they're all interactive. TV isn't—not yet, anyway.

As it happens, our family withdrew cold-turkey from TV this morning. We called Verizon and canceled our FiOS TV service. The set-top box, Linux innards and all, is now sitting in the

hall, waiting to be ferried back to a Verizon office.

The reason was choice. Even at its best, TV didn't give us much—not compared to the endless millions of choices on YouTube, Hulu and everybody else with video to share on the Web.

The free stuff—old-fashioned over-the-air (OTA) TV—is a mess. By the time you read this, most or all of the TV stations in the US will be transmitting digital audio and video, via ATSC. Old-fashioned analog NTSC, which has been with us since the 1940s, will be gone by the June 12 deadline. I'm not sure how much people will bother watching. All you get are a couple dozen channels, tops.

On cable or satellite, you can get much more. I don't think you can get a bigger selection than what Verizon FiOS offers. Where we live, FiOS carries 596 channels, including 108 HD channels and 136 premium channels, most of which are also HD. By the time I canceled the service on the phone this morning, the FiOS agent had reduced the price of the Extreme HD plan to \$47.99/month, including free DVR set-top box rental (normally \$12.99/month). That plan has 358 channels, including all 108 HD channels. It's a helluva deal, if you like a lot of TV. Making FiOS even better is that it comes over a fiber-optic connection that provides uncompromised data quality.

But we still canceled it, because we'd rather not watch channels at all. We'd rather watch programs. Or movies. Or stuff that doesn't fit either category. And, we'd prefer a better way to select them than by struggling with any of the cable or satellite systems' "guides", which are all terrible. It's much easier to navigate file paths and to do it on a real computing device,

including today's smart phones (which are really data devices that also do telephony). Because there's lots of video available on-line and from rental services like Netflix, we figure we'd take advantage of those. As it happens, Verizon makes it easy to get them in high-def, because we remain customers of FiOS high-speed Internet. There we get a solid 20Mbps both upstream and down, for \$64.99. It's an excellent deal, because that's for the whole world, and not just for a few hundred "channels" behind the gate to a walled garden.

Now that we've walked out of cable TV's walled garden, I can see how it traps the carriers even more than it traps the viewers. What they're trapped in is a scarcity game. And, they're losing. The producers and consumers are getting together without them. I can watch ACC sports on-line at the Raycom site. Nearly every channel on TV has a Web site that offers either live or archived content. True, all of them are pains in the butt to use (some requiring Flash plugins or worse), and many make half-hearted efforts to protect their cable and station distributors. But the writing is on the screen.

Now I'm thinking about what the abundance business would be like. What would you want out of the carriers if their Linux set-top boxes were open, or if you could provide your own? What game should they be playing once all they own is, say, the railroads and not the whole *Monopoly* board? Or hey, choose your own metaphor. Let's help them out here. They'll need it. ■

Doc Searls is Senior Editor of *Linux Journal*. He is also a fellow with the Berkman Center for Internet and Society at Harvard University and the Center for Information Technology and Society at UC Santa Barbara.

You know one hour of downtime is a big deal. Make sure your hosting provider does, too.



Your E-commerce concerns center around availability, security, compliance and scalability. Our responsibility as Linux hosting experts is to help you eliminate those worries when it comes to the hosting infrastructure behind your site. Bottom line — if your E-commerce site needs to be online and stay online, it needs to be hosted at Rackspace.

With our 200 RHCEs, we have the expertise and support to keep you online all of the time. And we understand the stakes — your reputation and revenue can suffer drastically from even a few minutes of downtime. You need to be up and running under any circumstances. **We work around the clock to make sure you are.**

Always Open



Microsoft

CISCO



AMD



rackspace.com/linuxjournal • 888-571-8976 | *experience fanatical support*

rackspace
HOSTING

More GFLOPS, Less WATTS

Intel® Nehalem is here!

**Higher Memory Bandwidth with DDR3 and QPI
Clusters and Servers Consume Less Power**

Four Servers in a 2U Chassis with all Hot-Swap:

- ▶ 1200 Watt 1+1 supply, 12 Drives, and Server Modules!

**FasTree™ ConnectX® QDR and DDR InfiniBand
Switches and HCAs**

Intel Professional Compiler Suite and Cluster Toolkit

- ▶ Version 11 with Nehalem Enhancements
- ▶ Academic Pricing Available



Configure your next Cluster today!

www.microway.com/quickquote



GPU Computing



WhisperStation™

With 1 to 4 Tesla GPUs

Tesla C1060 GPU Performance:

- ▶ 1 TFLOPS per GPU
- ▶ 4 GB DDR3 per GPU
- ▶ 102 GB/Sec Bandwidth
- ▶ CUDA SDK

Run MATLAB® on Tesla with "Jacket"

Clusters With Tesla™

S1070 - 4 GPU Servers

- ▶ 36 GPUs + 36 CPUs + 24 TB in 24U
- ▶ 40 Gbps FasTree™ InfiniBand
- ▶ InfiniScope™ Network Monitoring

FREE 15-day trial available
at microway.com

Microway
Technology you can count on™

508-746-7341
microway.com



GSA Schedule
Contract Number:
GS-35F-0431N