

jQuery

Squid

PXE

FreeBoo

Munin

Solaris-Zones

Irrlicht

LINUXTM JOURNAL

Since 1994: The Original Magazine of the Linux Community

APRIL 2009 | ISSUE 180

Run Linux and Solaris
on **One Machine**

Build an Unattended
Network-Based OS
Installer with **PXE**

Use **FreeBoo** to
Restore and Boot

Build a Secure
Squid Web Proxy

SYSTEM ADMINISTRATION

Program 3-D Graphics
with **Irrlicht**

Monitor Anything
with **Munin**

How to Bring Back a
Master Boot Record



PLUS POINT/COUNTERPOINT: **Mutt vs. Thunderbird**




The Embedded Wireless Dream Machine.

680MHz CPU, 64MB RAM, 16MB Flash, USB,
3 mini-PCI slots, 5A power supply for multiple
hi-power radios, and enhanced temperature
performance and ESD protection for carrier
applications. Full Linux SDK support and a
\$59 USD MSRP.

RouterStation



 www.ubnt.com

RouterStation UI Challenge: For Contest Details Visit: www.ubnt.com/challenge
Cash Prize \$200000 (Two hundred Thousand) USD.

User Interface development based on OpenWrt Linux firmware. OpenWrt is a registered trademark of OpenWrt.

CASH PRIZE

\$200000

www.ubnt.com/challenge

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

SOHO SERVER / STORAGE

OPTIMIZED FOR SMALL OFFICE & HOME OFFICE APPLICATIONS

SDP-S4100-T00

COST EFFECTIVE ENTRY LEVEL SERVER

- Mid Tower Chassis with 300W Power Supply
- Supermicro Server Board with Intel® 945GC Chipset
- 2x Realtek RTL8111C-GR Gigabit LAN Ports
- Support up to 4x SATA (3.0Gbps) Ports RAID 0, 1, 5, 10
- 4x 3.5" internal drive bays

STARTING AT

\$329


SDP-S4200-T00

MOST COST EFFECTIVE LAN SOLUTION

- Mid Tower Chassis with 300W Power Supply
- Supermicro Server Board with Intel® 3200 + ICH9R Chipset
- Support up to a Dual-Core Intel® Xeon® 3000 Series processor
- Support up to 6xSATA (3.0Gbps) Ports RAID 0, 1, 5, 10
- 4x 3.5" internal drive bays

STARTING AT

\$439



SDR-S1208-T00

STARTING AT **\$559**

- Supermicro Mini 1U Rackmount Server with 260W Power Supply
- Supermicro Server Board w/Intel® 946GZ Chipset
- Support up to a Dual-Core Intel® Xeon® 3000 Series processor
- TPM Support
- 1x 3.5" Internal Drive Bay
- 2x Intel® 82573 PCIe Gigabit LAN Port



SDR-C1303-T02

STARTING AT **\$899**

- 1U Rackmount Server with 400W Power Supply
- Supermicro Server Board w/ Intel® 5100 Chipset
- Support Dual Intel® 64-bit Xeon® Quad-Core or Dual-Core
- Support up to 48GB 667/533MHZ DDR2 ECC Reg
- 2x 3.5" Hot-swap SATA Drive Bay
- Intel® 82573V & 82573L Gigabit Ethernet Controller



SDR-C2301-T06

STARTING AT **\$999**

- 2U Rackmount Server with 460W Power Supply
- Supermicro Server Board w/ Intel® 5100 Chipset
- Support Dual Intel® 64-bit Xeon® Quad-Core or Dual-Core
- Support up to 48GB 667/533MHZ DDR2 ECC Reg
- 6x 3.5" Hot-swap SATA Drive Bay
- Intel® 82573V & 82573L Gigabit Ethernet Controller



SDR-C3301-T16

STARTING AT **\$1,359**

- 3U Rackmount Server with 650W Power Supply
- Supermicro Server Board w/ Intel® 5100 Chipset
- Dual Intel® 64-bit Xeon® Quad-Core or Dual-Core
- Support up to 48GB 667/533MHZ DDR2 ECC Reg
- 16 x 3.5" Hot-swap SATA Drives Trays
- Dual-port Gigabit Ethernet Controller



SDR-C4302-T02

STARTING AT **\$899**

- 4U Rackmount Server with 600W Power Supply
- Supermicro Server Board w/ Intel® 5100 Chipset
- Support Dual Intel® 64-bit Xeon® Quad-Core or Dual-Core
- Support up to 48GB 667/533MHZ DDR2 ECC Reg
- 3 x 3.5" Internal SATA Drives Trays
- Dual-port Gigabit Ethernet Controller



SDR-A8301-T42

STARTING AT **\$4,799**


- 8U Chassis with 1350W Redundant Power Supply
- Supermicro Server Board w/ Intel® 5400 Chipset
- Support Dual Intel® 64-bit Xeon® Quad-Core or Dual-Core
- Support up to 64GB SDRAM Fully Buffered DIMM (FB-DIMM)
- 42 x 3.5" Internal SATA Drives Trays
- 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 assemble systems that require a little extra. Servers Direct - your direct source for scalable, cost effective server solutions.

1.877.727.7887 | www.ServersDirect.com

Intel, Intel logo, Intel Inside, 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

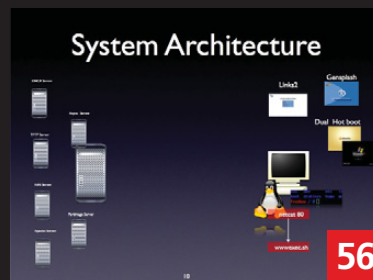

APRIL 2009

Issue 180

SYSTEM ADMINISTRATION

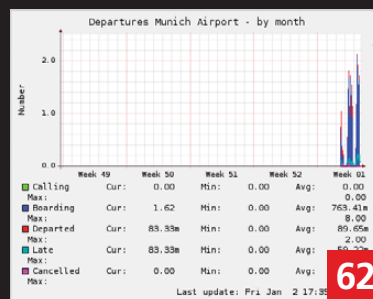
FEATURES

- 50** **PXE: NOT JUST FOR SERVER NETWORKS ANYMORE!**
Build a Windows install server...using Linux!
Bill Childers
- 56** **FREEBOO: AN OPEN ARCHITECTURE FOR NETWORK DUAL BOOT**
Open-source network restore and boot.
Cristina Barrado and Sebastian Galiano
- 62** **MUNIN—THE RAVEN REPORTS**
Monitor anything, not just system parameters.
Patricia Jung



ON THE COVER

- Run Linux and Solaris on One Machine, p. 68
- Build an Unattended Network-Based OS Installer with PXE, p. 50
- Use FreeBoo to Restore and Boot, p. 56
- Build a Secure Squid Web Proxy, p. 36
- Program 3-D Graphics with Irrlicht, p. 73
- Monitor Anything with Munin, p. 62
- How to Bring Back a Master Boot Record, p. 40
- Point/Counterpoint: Mutt vs. Thunderbird, p. 77





Systems



ZT Systems delivers something different: a unique

BALANCE

of world-class server performance and cost advantage joined with extensive flexibility in delivery and support

From the small office to the mega-datacenter, ZT Systems provides dependable, energy efficient server and client solutions that maximize productivity and optimize your IT budget. An established manufacturer celebrating our 15th anniversary, ZT combines individualized attention to your requirements with the experience and scalability you demand. All ZT Systems are assembled and serviced in the USA.



Choose ZT Systems and get:

- Custom Configurations
- Professional Grade Quality
- Personalized Account Service
- Superior Value



ZT Reliant SC1256i-40 Small Business Server
Entry-level Single Socket Solution

Starting at Only \$749

ZT-SC1256i-40

- Intel® Core™ 2 Quad Processor Q8200
- 4GB DDR2 800MHz SDRAM (2 x 2GB)
- (2) 1.5TB SATAII 32MB Hotswap Hard Drives (Max. 4)
- 16X DVD-ROM Drive
- RedHat Enterprise 5 Trial Version
- 3 Year Limited Warranty and 24x7 Telephone support



ZT Systems 1105Ri-82 1U Rack Server
Entry-level file/web/email Server Solution

Starting at Only \$949

ZT-1105Ri-82-C000001

- Intel® Xeon® Processor X3220
- 4GB DDR2 800MHz ECC Unbuffered SDRAM (2 x 2GB)
- (2) 500GB SATA Enterprise Hard Drives (Max. 4)
- 16X DVD-ROM Drive
- RedHat Enterprise 5 Trial Version
- 3 Year Limited Warranty and 24x7 Telephone support



ZT Systems 1213Ri-84 1U Rack Server
Data Center / Network Infrastructure and Cluster Solution

Starting at Only \$1249

ZT-1213Ri-84-C000001

- Single Intel® Xeon® Processor E5410 (dual capable)
- 4GB DDR2 ECC/REG. SDRAM (2 x 2GB)
- (2) 500GB SATAII 32MB Hotswap Hard Drives (Max. 4)
- Slim DVD-ROM (Optional)
- RedHat Enterprise 5 Trial Version
- 3 Year Limited Warranty and 24x7 Telephone support



ZT Systems 4202Ci-84 4U/Tower Rack Server
Convertible Tower / Rack SMB Solution

Starting at Only \$1649

ZT-4202Ci-84-C000001

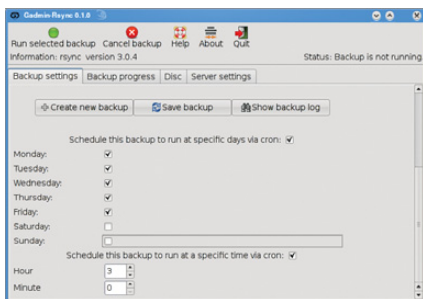
- Single Intel® Xeon® Processor X5410 (dual capable)
- 8GB DDR2 ECC/REG. SDRAM (4 x 2GB, Dual Rank)
- (3) 1TB SATAII 32MB Hotswap Hard Drives (Max. 8)
- 16x DVD-ROM & 1.44MB Floppy Drive
- RedHat Enterprise 5 Trial Version
- 3 Year Limited Warranty and 24x7 Telephone support

Scalable Custom Server Solutions - Contact Us to Learn More
(866) 984-7687 corpsales@ztsystems.com

CONTENTS APRIL 2009 Issue 180

COLUMNS


- 10 SHAWN POWERS'**
CURRENT_ISSUE.TAR.GZ
Fixing Starships, Transporters and Even Computers
- 22 REUVEN M. LERNER'S**
AT THE FORGE
Writing jQuery Plugins
- 26 MARCEL GAGNÉ'S**
COOKING WITH LINUX
Backing Up to the Clouds



- 32 DAVE TAYLOR'S**
WORK THE SHELL
More Fun with Word and Letter Counts
- 36 MICK BAUER'S**
PARANOID PENGUIN
Building a Secure Squid Web Proxy, Part I
- 40 KYLE RANKIN'S**
HACK AND /
When Disaster Strikes: Scratching the Master Boot Record
- 77 KYLE RANKIN AND**
BILL CHILDER'S
POINT/COUNTERPOINT
Mutt vs. Thunderbird
- 80 DOC SEARLS'**
EOF
Exchanging Exchange

INDEPTH

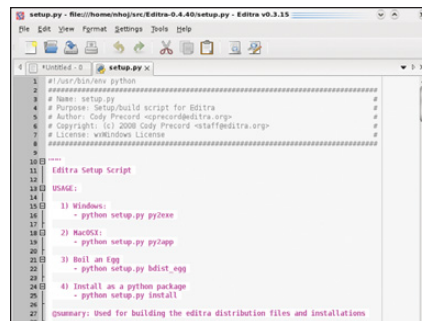
- 68 SOLARIS-ZONES: LINUX IT**
MARBLES GET A NEW BAG
Run Linux and Solaris on the same machine.
Victor Burns
- 73 3-D GRAPHICS**
PROGRAMMING WITH
IRRLICHT
3-D graphics programming—not just for geniuses anymore!
Mike Diehl



IN EVERY ISSUE

- 10 FROM THE EDITOR**
- 12 LETTERS**
- 16 UPFRONT**
- 44 NEW PRODUCTS**
- 46 NEW PROJECTS**
- 65 ADVERTISERS INDEX**

44



46 NEW PROJECTS: EDITRA

Next Month

COOL PROJECTS

Here at *LJ*, we hardly even get excited anymore when somebody mentions a new project that uses Linux, but that's not the case with these Cool Projects. We've got Linux running an un-manned submarine and rocket, and if that doesn't get your BP up, you may already be dead.

If it's all too much for you, don't worry; we've got articles for the weak heart types also—articles on setting up home control and home surveillance systems. Plus, one on converting your old 8mm movies so that you can sit back and reminisce about the days when you weren't almost dead.

If you still have doubts about the ultimate success of Linux and open source, don't miss our upcoming interview with Neuros Technology CEO and founder Joe Born. Find out how open source concepts transformed his company. If his statements don't convince you, you definitely are already dead



Experience Lightning Without The ThunderSM



▶ **WhisperStation-ProTM**

1 TERAFL0P IN A COOL, FAST, RELIABLE PLATFORM!

Whether it's Wall Street, Main Street or Your Street, Microway's new Nvidia-powered WhisperStation-Pro is energy-efficient, designed for superior performance, and best of all - QUIET.

Originally designed for a group of power hungry, demanding engineers in the automotive industry, WhisperStation-Pro incorporates two AMD[®] Opteron[™] or Intel[®] Xeon[®] quad-core processors and high-efficiency power supplies. Ultra-quiet fans and internal sound-proofing produce a powerful, but silent, computational platform.

WhisperStation-Pro configured with one Quad core processor, 4 GB high speed memory, 250 GB drive, dual-GigE, NVIDIA[®] Quadro[™] FX570 graphics and 20" LCD – **starts at \$1995.**

You can have it configured to your exact needs with NVIDIA GeForce[®] or Quadro graphics adapters (including SLI[®]), NVIDIA Tesla[™] GPU, any Linux distribution, or Windows[®] dual-boot. Also, there is plenty of room for RAID storage expansion. From a home based workstation for financial wizards, to a superior gaming or design station, WhisperStation-Pro fits the bill and your budget.

Visit www.microway.com for more technical information.

Hear Yourself Think Again!

Call our technical sales team at 508-746-7341 and customize your WhisperStation-Pro today.

**WhisperStation[™] ▶
3D Elite SLI
For Gamers**



Microway[®]
Technology you can count on[™]

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
mggagne@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.



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

Everything you need for

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

50%
OFF



a successful website

DOMAIN NAMES

.biz

Register your business website and save!
Protect your personal information – private domain registration is included for FREE!

**.biz domains ~~\$8.99/year~~
now \$3.99/first year***



WEB HOSTING



FREE Domain Names!

The 1&1 Business Package gives you everything you need for a professional website. 3 domains, site-building tool, 2,500 e-mail accounts, search advertising vouchers and more!

**Starting at ~~\$9.99/month~~
now \$4.99/mo.* (first 6 months)**



PREMIUM SERVERS



Designed specifically for high performance needs, these top-of-the-line AMD™ processors feature 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.

**Starting at ~~\$199.99/month~~
now \$99.99/mo.* (first 6 months)**



E-COMMERCE SOLUTIONS



Set up your online store and start selling. Integration with eBay®, Shopzilla™, Google Product Search™ and Shopping.com® included!

**Starting at ~~\$24.99/month~~
now \$11.99/mo.* (first 6 months)**



*Offers valid for a limited time only. 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

Fixing Starships, Transporters and Even Computers

As a system administrator, it reassures me when I see the computer systems break in *Star Trek* episodes. The reason it brings me peace of mind is that I know if I perfect human cloning and mind transfer, my sysadmin skills still will be useful in a few thousand years. Also, I'm fairly certain Data (the android Lieutenant) must be running some sort of embedded Linux, so my particular skill set will be very much in demand.

I'm getting a little ahead of myself, however. Although I'm sure system administration will be a popular topic in future millennia, with this issue of *Linux Journal*, you hold the future in your hands right now. Cheesy anecdotes aside, this month we have tons of articles and columns to help you keep the packets flowing.

Because most system administration begins with installing an operating system, it seems appropriate to start there. We've got a couple really interesting articles this month that deal with both system installing and system restoring. Bill Childers tells us how to leverage PXE network booting to install operating systems (even Windows!) by using a remote booted Linux kernel. Once you understand the nuances of PXE booting, it's amazing the things you can do without even a local boot device. Add to that Christina Barrado and Sebastian Galiano's article on FreeBoo, and you'll be able to install, restore and dual boot in no time (or at least, in less time than with conventional methods). FreeBoo is an open-source alternative to programs like Rembo. Using PXE and FreeBoo will make booting and restoring multiple operating systems much easier than ever before.

Once your systems are up and running, it's the system administrator's job to keep them that way. Applications such as Munin can keep track of long-term trends. In fact, Munin can keep track of more than just system logs, and Patricia Jung tells us all about it. Whether you use Munin, or you just grep system logs, when the inevitable problem comes along, it's important to know what to do. And, Kyle Rankin explains how to salvage a hard drive when the mysterious Master Boot Record fails. Thankfully, he goes into a little more detail than the standard Windows answer. Sometimes, `fdisk /mbr` isn't enough. GRUB is a bit more robust than that, thank

you very much.

Victor Burns shows us an intriguing new method for dual booting between Linux and Solaris. No, I'm not talking about full virtualization or standard dual booting, I'm talking about running both at the same time. With Solaris-Zones, you really can have the best of both worlds. Speaking of both worlds, Bill and Kyle hash it out again this month. This time, it's with e-mail clients. I won't say Bill is GUI, but he sure prefers it in an e-mail client.

If you add our regular cast of columnists to this issue, you'll see it's quite a great month to be a *Linux Journal* subscriber. Mick Bauer shows us how to set up the Squid Web proxy securely. Dave Taylor hones our scripting skills, and Marcel Gagné shows us one of the most important things a system administrator can do—back up. If you're not a sysadmin, don't worry about being left out this month. Even though we all need to administer our own systems to some degree, we've also got information that will appeal to those readers without racks of servers to manage.

Mike Diehl teaches us how to program using Irrlicht to get some awesome 3-D graphics in our programs. Reuven M. Lerner shows us how to write plugins for jQuery. Plus, we have tech tips, new product information and news from the Linux industry that is bound to tickle the fancy of any Penguin fan.

Perhaps none of us might be around when it's time to do tech support for Galaxy-class starships. It's unlikely we'll be able to debug code for a holodeck any time soon. And, to be honest, I don't expect to get a call requesting my help in programming any interstellar guidance systems. I think it's a fair guess, however, that when the time does come, most of the devices will be running the Linux kernel. If we practice now, perhaps we'll be ready if the time comes sooner than I expect. Or, if we get visited in a time-travel episode. Admittedly though, the latter is probably unlikely. ■

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://www.freenode.net) IRC channel on Freenode.net.

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, 0, 1, 10
- Hot Swap, Hot Spare
- Linux, Windows, Mac
- E-mail Notification
- Tower or Rackmount

The new generation of energy-efficient 45W AMD Athlon™ X2 dual-core processors with Cool'n'Quiet™ technology are designed to enable quiet performance.



Silent Eco Green PC

Based on the new 45W AMD® Athlon™ X2 dual-core processor is energy efficient, quiet and has plenty of power. starts at **\$199**

Fanless Silent ITX PC

1G DDR2, 80GB Hard Drive starts at **\$299**
AMD® processor, Low-profile Add-on Available




4U 24Bay Storage Server

AMD Opteron™ Processors
36TB RAID-6 128GB RAM, 4 x GigaLAN



Small 1U Server for Data Center ISP

AMD Phenom™ processor
4 to 8GB DDR2-800, 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

AMD, AMD Athlon, Phenom, Opteron, Cool'n'Quiet, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other names are for informational purposes only and may be trademarks of their respective owners.



Where's the Security?

I was surprised in a *Linux Journal* issue (January 2009) specifically targeting security that the shell script article by Dave Taylor didn't suggest a more secure method of creating temporary filenames other than just using the \$\$ variable. Adding a little `${RANDOM}` to the temp directory name and filename, perhaps? Maybe he's going to address that in a later article.

--
Forrest Hudspeth

Dave Taylor replies: *Great point, Forrest. Since I have limited space in my column, I tend to write short solutions rather than more complicated ones. You're right that something like:*

```
tempname="/tmp/appname.$$.${RANDOM}"
```

works better than just a simple application of \$\$ in terms of making it harder for a hacker to game your temporary files and potentially alter the execution and results of system shell scripts.

Lots of Security

In my opinion, the January 2009 Security issue was one of the best issues put out in a while, not that the magazine isn't good, but you really outdid yourselves this time in keeping the articles on focus

with the focus of the issue. The one thing I would have liked to see is if Cooking with Linux could have talked about setting up home surveillance cameras in keeping with the Security theme or something of the like. Although I always enjoy Marcel Gagné's articles, I'd like to see him focus on some small-to-medium projects too instead of just games.

--
Mike

Laptop Sizes and Great LJ Subscription Service

In the January 2009 issue, you have a nice article: "Small Laptops vs. Large Laptops" by Kyle Rankin and Bill Childers. Well, I've worked for several years as a consultant, and I travel a lot, so I can understand the need for a big laptop and also can understand the need and usefulness of a small laptop. But, why not have the best of both worlds? First, I had a big 15" or 17" laptop, and it was great, but then I had to travel, and it wasn't so nice to carry nor was it usable on an airplane, so I bought a small Sony with an 11" monitor. It was nice to have eight hours of battery life, but at the end of the day I was tired, and my head hurt because of the effort to read all the tiny letters. The keyboard was too small, and I spent much time using delete and backspace to correct my errors while programming or writing documentation. So, I bought a 13" display laptop, once again from Sony—an S2 series. Right now, I'm the happiest man in the world with my laptop. It's light (1.5kg), fast, has 2.0MHz dual core and 4GB of RAM. It has a nice hybrid dual video card (one Intel for extreme endurance on battery, and one NVIDIA for games and so on). Nowadays, I think we can have both worlds, being light and small and at the same time, powerful and functional.

On another subject, some weeks ago I had a problem with my subscription for *Linux Journal*, and I didn't get my favorite magazine. I contacted *LJ* support and they were great. In a week, I had the missing issues (and I live in Portugal). Here's something we all love:

real support when you need it, and as with Linux, *LJ* gives you the best support. Keep up the excellent work!

--
Pedro

Pedro, glad you liked Kyle and Bill's new Point/Counterpoint column—see page 77 of this issue for more bickering.—Ed.

Reducing Keystrokes

Further to Antoine's suggestion to Dave Taylor on reducing keystrokes [see Letters in the January 2009 issue of *LJ*], note that:

```
pickline=$(( $RANDOM % 250 + 1 ))
```

can be further simplified to be:

```
pickline=$(( RANDOM % 250 + 1 ))
```

It seems that all bash environment variables are expanded within the scope of `$(())`.

--
Black Jack Shellac

More Love for Kyle and Bill

Just wanted to drop y'all a line to voice my appreciation for Point/Counterpoint. While Kyle vs. Bill did smack of a rap battle in a few places, it was excellent nonetheless. Keep up the great work, guys!

--
Samuel

Ouch!

The "Go Green, Save Green with Linux" article [by James Gray, in the April 2008 issue of *LJ*] carelessly misquotes an incompetently uttered sound bite of Springboard Research: "an average-sized server has the same carbon footprint as a mid-sized 4wd taking 17 litres of fuel to travel 100km", said Bob Hayward, Research Vice President for Springboard Research. His comment says nothing. Is he comparing a server to 17 liters of fuel? or a 4wd? Or 100km, walking, flying, driving? *Linux Journal* butchered this incomprehensible statement into the ludicrous, saying a server has the same footprint as a 4WD vehicle! This is business incompetence



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—technology community.

This year's JavaOne conference offers even more opportunity to grow your language 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 JavaOne conference Rock Stars and Java technology creators and evangelists
- Experience JavaFX™, RIAs, and other 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!

Save \$200
on Conference registration!
Register by April 22 at
java.sun.com/javaone

Cosponsors



Media Cosponsors



Media Partners

Application Development Trends

Extension Media

Linux Journal
Slashdot.org

SourceForge.net

Toolbox.com



© 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.

[LETTERS]

stacked on journalism incompetence.

--
Jim Leuba

James Gray replies: I personally find Hayward's quote ingenious, which does two things. First, Hayward compares the energy consumption of two very different machines: one whose energy consumption isn't readily apparent (the server) to one that is more publicized (the vehicle). Second, he shows that servers consume more energy than their compact size belies. Hayward's metric of energy consumption is the carbon footprint, which is simply the quantity of carbon dioxide created by the burning of fossil fuels in energy generation—pretty straightforward stuff.

I can only speculate why you had such a strong emotional reaction to Hayward's quote. My best guess is because it is hard for us human beings to accept that

our actions are directly and gravely degrading the health of our planet. It's tough news to hear.

Tip

This is in reference to David Sinck's Tech Tip in the December 2008 issue. I find the following code a little simpler and also insensitive to file extensions:

```
#!/bin/sh
case `file -b $1 | cut -d' ' -f1` in
  Zip|gzip) CAT=zcat ;;
  bzip2) CAT=bzcat ;;
  *) CAT=cat ;;
esac
$CAT $1
```

As an aside, if one is looking to open the files in an editor, vim makes it easy by opening zip files in various formats.

--
Mayuresh Warunjikar

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.



Photo submitted by Brenton Leanhardt: "I made this when I was in Colorado for Christmas."

LINUX JOURNAL

At Your Service

MAGAZINE

PRINT SUBSCRIPTIONS: Renewing your subscription, changing your address, paying your invoice, viewing your account details or other subscription inquiries can instantly be done on-line, www.linuxjournal.com/subs. Alternatively, within the U.S. and Canada, you may call us toll-free 1-888-66-LINUX (54689), or internationally +1-818-487-2089. 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.

CommunityOne

An open developer conference

june 1-3
moscone center
san francisco

MySQL™

Groovy

Cloud

Web 2.0

Java™

NetBeans™

RIA

Apache

Virtualization

Eclipse

Linux

PHP

Android

OpenSolaris™

Open

GlassFish™

Free*

**“one
heavyweight
geek-fest”**

—Tim Bray, 2008 speaker

OK, the economy is bad, but your development choices don't have to follow suit. Whether you're focused on creating robust Web apps, building a scalable infrastructure, or thinking about cloud computing, you'll find sessions on proven free and open-source software (FOSS) technologies and tools. This is stuff you can really get excited about.

Register Today!

*Monday is free; fee required for tutorials on Tuesday and Wednesday.

developers.sun.com/events/communityone

© 2009 Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, GlassFish, Java, MySQL, NetBeans, and OpenSolaris 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.

CommunityOne, 999 Skyway Road, Suite 300, San Carlos, CA 94070



diff -u

WHAT'S NEW IN KERNEL DEVELOPMENT

The **kernel boot process** is one of the strangest parts of the whole operating system. One aspect of it is **initramfs**, where the kernel must first boot up in a RAM-based root filesystem, before it can mount the disk-based filesystems and transition the kernel to consider those the root filesystem. **initramfs** traditionally has not been part of the official kernel, so each Linux distribution has come to maintain its own unique implementation. In general, this means the sequence of events in the boot process of each distribution is different. This can affect the way hardware is initialized and the state it's in when the user finally sees a login prompt.

Jeremy Katz recently announced **Dracut**, a new distribution-independent implementation of **initramfs**, inspired by discussions at the most recent **Kernel Summit**. The project's goal is to simplify and standardize the boot process and move the **initramfs** implementation into the official kernel tree. Each Linux distribution then would boot up identically, but would be able to layer any additional functionality on top of the main **initramfs** implementation.

Jeremy's work is based on the **Red Hat initramfs**. **Hannes Reinecke** is working on his own implementation, based on **OpenSUSE**. It's likely the two projects will merge at some point. The project will need to bear several issues in mind. As **Neil Horman** explained, some systems elect to stay in the RAM-based filesystem for stability reasons and never transition to a disk-based root filesystem. Any successful **initramfs** would have to support that feature, in order to work with various embedded system projects.

Keeping track of what goes into the kernel is

always pretty tough. Each new version could contain thousands of patches. There are various ways to track them all, including using **git** or one of the git repository browsers. **John Hawley** recently set up a new method at **patchwork.kernel.org**. This page keeps track of all patches sent to a variety of kernel-related mailing lists, including the main one, **linux-kernel**. It presents the patches in color and is nicely formatted, and it lets folks change the status of each patch, so it's clear which have been accepted.

Although a variety of git documentation exists, some of it is fairly involved. For folks who'd rather just learn what they need to know, **David Howells** has written "The Git Hater's Guide To The Galaxy" and submitted it for inclusion in the kernel source tree. It's not the most in-depth git doc out there, but it covers a lot of the stuff David himself uses it for on a daily basis, and it may be more accessible than the more detailed docs out there.

One problem with running **dæmons** is that the **init** process loses track of the relationship between the **dæmon** process and the parent processes that spawned it. As **Scott James Remnant** has discovered, there doesn't seem to be any straightforward way in Linux for **init** to know by which process a given **dæmon** was spawned.

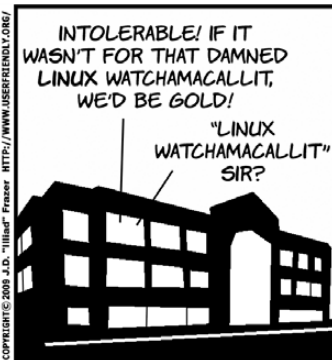
That information is currently lost by the spawning procedure.

Scott has coded up support for some new signals to carry this information into the kernel. If accepted into the main tree, this code will allow system logs to track effectively the activities of the application that spawns any **dæmons**, along with the **dæmons** it spawns. Among other things, this will improve debugging efforts for any applications that spawn **dæmons**.

Most modern CPUs have special registers for tracking performance. Things like **perfmon** use these, but recently **Thomas Gleixner** and **Ingo Molnar** took a shot at designing something better. Their initial effort actually has met with fierce resistance from **Paul Mackerras** and **David S. Miller**, who felt that the approach Thomas and Ingo had taken was not likely to meet the needs of real systems. The debate does not seem anywhere near over, and certainly Thomas and Ingo will produce some code to prove their points. Probably something excellent will come out of all this, but currently, the two sides have not had much of a meeting of the minds. Paul has said there's no way to track a process effectively without stopping it, and Thomas and Ingo say, yes, it can be done. Time will tell.

—ZACK BROWN

USER FRIENDLY by J.D. "Iliad" Frazer



LINUX JOURNAL EDITION

1. Lowest legal price in dollars of the *Ghosts I–IV* album by Nine Inch Nails: **0**
2. Price in dollars charged for *Ghosts I–IV* at Amazon.com: **5**
3. Position of *Ghosts I–IV* among Amazon's top-selling albums for 2008: **1**
4. Millions of female drivers in the UK who are "putting themselves and other drivers at risk by wearing the wrong footwear when behind the wheel": **11.5**
5. Preference percentage for a "higher resolution" photo that isn't, in a choice between two actually identical pictures: **75**
6. Billions of dollars spent annually on Internet advertising by 2011: **106**
7. Percentage of US Internet audience that viewed on-line video in November 2008: **77**
8. Billions of on-line videos viewed by US Internet users: **12.7**
9. Percentage increase of above over the same period in 2007: **34**
10. Google's (Linux-served) YouTube rank as a US on-line video property: **1**
11. Billions of YouTube videos viewed in November 2008: **5.1**
12. Google Sites again ranked the top US video property, with nearly **5.1 billion videos**.
13. YouTube percentage of all videos viewed in November 2008: **40**
14. YouTube share of Google video properties: **98**
15. Millions of US Internet viewers in November 2008: **147**
16. Average number of videos users watched by US viewers in November 2008: **87**
17. Duration in minutes of the average on-line video: **3.1**
18. Duration in minutes of the average on-line video viewed at (Linux-served, via Akamai) Hulu.com: **11.9**
19. Number of results in a search for "Linux" at AMD.com (US): **1,281**
20. Number of results in a search for "Linux" at Freescale.com: **1,445**

Sources: 1: *Nine Inch Nails* | 2, 3: *Amazon.com* | 4: *Sheila's Wheels* (a UK insurance company), citing a 2007 YouGov survey of 754 female drivers | 5: *ArsTechnica*, citing *Journal of Consumer Research* | 6: *CRM Today* | 7–9 and 11–17: *MarketingVOX*, sourcing comScore's *Video Metrix, Marketing Charts* | 10, 18: *SearchDNS.Netcraft.com* | 19: *AMD* | 20: *Freescale*

EUCALYPTUS: a Tree Growing in the Cloud

From the Linux (and *Linux Journal*) perspective, there's an issue with clouds—those back-end Web services that compose Utility Computing. They're proprietary. Amazon owns AWS (Amazon Web Services: S3, EC2 and a growing number of others). Google, Microsoft and other companies own theirs as well.

These are open in the sense that most services are substitutable, which makes them utilities. If you don't like storing your data at Amazon's S3, you can park it elsewhere. The documents you create and keep at Google Docs can be put elsewhere as well. Elastic computing is more tricky, but it should be substitutable as well.

In a way, all these data centers are available to the public as a utility service. They are an important trend in computing, comprising the growth of a vast open SLA (service-level agreement) infrastructure for computing resources spread across the Net. (Nicholas G. Carr wrote about all this in his book *The Big Switch—Rewiring the World, from Edison to Google*. I visited the topic in "The Bigger Switch", my EOF column in the May 2008 issue of *Linux Journal*.)

But, why should The Cloud consist only of proprietary back ends? How about building our own cloud services, our own way?



Eucalyptus

I was pleased to find an answer from colleagues at UCSB (where I'm a fellow at the Center for Information Technology & Society), namely Rich Wolski and his merry band of computer scientists in the Computer Science Department. There they have created EUCALYPTUS (Elastic Utility Computing Architecture for Linking Your Programs To Useful Systems). It's what they call "an open-source software infrastructure for implementing Elastic/Utility/Cloud computing using computing clusters and/or workstation farms". It is interface-compatible with Amazon's EC2 (the de facto standard for elastic cloud computing) and designed for modification and extension as needed for you-name-it client-side interfaces. Most important for *Linux Journal* readers, "EUCALYPTUS is implemented using commonly available Linux tools and basic Web service technology making it easy to install and maintain", they say.

The issue for IT folks is a range of choices that should include DIY as well as commercial back ends. Rich Wolski puts it this way:

Cloud computing permits "self-service" IT independently of whether the provider is an external company or an internal IT organization. Part of what makes IT onerous these days is that the infrastructure organization also must be a customer service organization. Cloud computing really streamlines customer service for IT in the same way that Web-based e-commerce streamlined customer service for retail. It doesn't eliminate the need entirely (you still can get someone on the phone even when a company has an e-commerce site), but it does improve the business efficiency considerably.

As I write this, version 1.5 is coming along. (Version 1.0 was released in June 2008.) And, there are ambitious plans. For more information, check out eucalyptus.cs.ucsb.edu.

—DOC SEARLS

SHOPPING ON PENGUINS



I was pointed recently to Zappos (www.zappos.com) as a near-perfect example of a company that brings the principles of open source to business. Its site is inventive and fancy (as you'd expect a clothing retailer to be), but not a triumph of design over utility. What's more, it's fast. I can check Web site responsiveness with some confidence, because our little apartment near Boston has 20Mb/s symmetrical service from Verizon FiOS (that's fiber, and pretty cheap, considering), and the speeds I get at the office I share at Harvard are more than twice that.

So I decided to see what they were running that site on, by checking with Netcraft.com's "What's that site running?" service. The answer was Linux. Zappos itself is in an Akamai Netblock, and of the 13 other results (all foo.zappos.com), the results for OS were Linux or "unknown".

Then I decided to look beyond Zappos to other on-line retailers. Here are the results for the top ten, as listed by InternetRetailer.com, with sales volume numbers in parentheses. Results for companyname.com are first, and other results for each company are summarized in text (Netcraft gives results in lowercase):

1. Amazon.com Inc. (\$14.8 billion): linux.
2. Staples Inc. (\$5.6 billion): linux, through akamai.
3. Office Depot Inc. (\$4.9 billion): the top result is linux for www.officedepot.com, but officedepot.com (without the dubs) is windows server 2003. Other foo.officedepot.com sites are a mix of the two and "unknown".
4. Dell Inc. (\$4.2 billion): f5-big-ip, though l.dell.com (Dell laptops) is linux and a couple of foo.dell.com sites are "unknown".
5. HP Home & Home Office Store (\$3.4 billion): hp-ux.
6. OfficeMax Inc. (\$3.2 billion): f5-big-ip, with a mix of solaris, windows server 2003, linux and "unknown" among other officemax sites.
7. Apple Inc. (\$2.7 billion): "unknown", but the company uses linux through akamai for a number of foo.apple.com sites.
8. Sears Holding Corp. (\$2.6 billion, includes Sears.com and Kmart.com): for sears.com, linux, plus some "unknown", solaris, f5-big-ip and windows server 2003 for sears.com.mx. For kmart.com, linux (through akamai), plus linux (with and without akamai) for various foo.sears.com sites. Kmart.com.au is windows 2000. Most of the rest are linux, plus a couple "unknown" and one f5-big-ip.
9. CDW Corp. (\$2.4 billion): the top result is linux for www.cdw.com (through akamai). cdw.com (sans dubs) is f5-big-ip. The rest are a mix of linux and f5-big-ip, with one windows server 2003.
10. Newegg: windows server 2003, though linux shows up in five out of the nine foo.newegg.com results.

Although that summation is far from a complete picture, or even a completely accurate one for this Top Ten, it's clear that the picture can't be painted without linux as the most primary color.

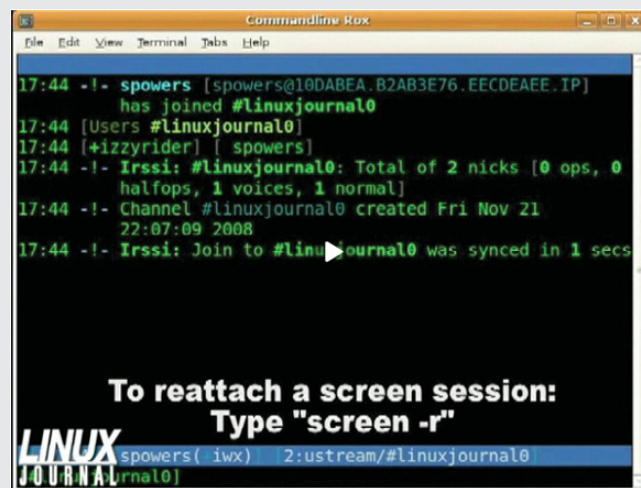
RESOURCES:

Internet Retailer Top 500 Retail Web Sites:
www.internetretailer.com/top500/list.asp
Netcraft: netcraft.com

—DOC SEARLS

Even More Videos on LinuxJournal.com

We love Web videos. We especially love videos that teach us new quick-and-easy tricks, so we've been cooking up some more videos over in the LinuxJournal.com kitchen. We'll have several short video tips per week in the Tech Tip of the Day section. Visit us often to see some fun tricks to make your life just a little better. Here's one of my favorites:
www.linuxjournal.com/video/transfer-your-terminal-screen.



—KATHERINE DRUCKMAN

LINUS' BLOG

That's what it's called. I found out about it when a pal sent a link under the subject line "Hell hath frozen over." Look up the above on Google and you'll find it. Blogspot, Netcraft tells us, is in the Google, Inc. netblock, and running on Linux.

—DOC SEARLS

They Said It

If computers of the kind I have advocated become the computers of the future, then computation may someday be organized as a public utility. We can envisage computing service companies whose subscribers are connected to them by telephone lines....

The system could develop commercially in fairly interesting ways. Certain subscribers might offer services to other subscribers. One example is weather prediction. A weather-predicting company that is a subscriber to the central computer predicts the weather but keeps the predictions in its private files. If you subscribe to its service, your programs can gain access to those files. You may even have weather-predicting programs run for your benefit to answer your own particular questions....Other subscribers might furnish economic predictions. The computing utility could become the basis for a new and important industry.

—**John McCarthy, in 1961**

At present, a newspaper, magazine or book is a package produced by a large organization. With the information utility, the physical production and distribution disappears, allowing a much smaller organization to put out the same packages of text and pictures. Moreover, the user does not face a one-shot decision to buy *Time* or *Newsweek*. He will be able to read the “cover” or table of contents of each, read such items as strike his fancy, and the system will bill him for what he reads from each source. In fact, since the cost of keeping a file of information in the computer and making it publicly available will be small, even a high school student could compete with the *New Yorker* if he could write well enough and if word of mouth and mention by reviewers brought him to public attention.

—**John McCarthy in the 1970s,**

www.theregister.co.uk/2005/11/10/gates_ozzie_memos/print.html and www-formal.stanford.edu/jmc/future/hoter1.html

Does anyone really think that a company-owned platform is going to win here, that it won't be swamped by an open federated system of servers that peer, like e-mail? If so, I'd like to hear why. We went through this exercise repeatedly in the tech industry; the lesson of history is clear—closed systems have their place and time, at the beginning of a new layer, when users need simplicity over everything else, they serve as training wheels when everyone is a newbie. Eventually we grow out of the need to have our hands held and the freedom of open systems becomes attractive, and we jump.

—**Dave Winer,**

www.scripting.com/stories/2009/01/03/helpingFriendfeed.html

Nobody will doubt the value of openness for the mobile industry anymore.

—**Sean Moss-Pultz,**

article.gmane.org/gmane.comp.handhelds.openmoko.community/33625

as a thank you to our fans for your continued support, we are giving away the new nine inch nails album one hundred percent free, exclusively via nin.com.

the music is available in a variety of formats including high-quality MP3, FLAC or M4A lossless at CD quality and even higher-than-CD quality 24/96 WAVE. your link will include all options—all free. all downloads include a PDF with artwork and credits.

for those of you interested in physical products, fear not, we plan to make a version of this release available on CD and vinyl in july. details coming soon.

—**Nine Inch Nails, dl.nin.com/thesip**

While Linux-based netbooks have not gained much consumer acceptance in the US, their success varies greatly by geographical market. In developing countries, Linux-based Eee PCs have fared better...Even if a majority of netbooks run Windows, the minority that run Linux are the most successful non-Windows, non-Macintosh consumer PCs in the industry in terms of penetration.

—**Forrester Analyst J P Gownder,**

www.vnunet.com/vnunet/news/2233381/netbooks-sales-driving-linux

FREESCALE

Freescale may be the first semiconductor company to associate itself aggressively with portable Linux devices. The former Motorola semiconductor division is sharply targeting the low-priced Linux-based Netbook market, which is hot in the world market and just starting to get warm in the US.

Early this year, Freescale launched a new processor called the i.MX515 Cortex A8 Netbook Processor. Company literature positions the part for “low-power, gigahertz performance Netbooks at sub-\$200 price points”. According to CNET, Glen Burchers, director of global marketing for Freescale’s consumer products group, says “Because the primary function (of a Netbook) is accessing the Internet, Linux and Firefox are a good operating system and application for that purpose.”

Freescale is developing a reference design with ASUS-subsidiary Pegatron. It will feature the 1GHz i.MX51 processor, Ubuntu Linux, Adobe’s Flash Player, a new power management chip and the SGTL5000 ultra-low-power audio codec.

Some data:

- The i.MX515 is a 65nm CPU based on an ARM11 Cortex-A8 blueprint.
- Clocks from 600MHz to 1GHz.
- Includes the SGTL5000 ultra-low-power audio codec.
- OpenVG and OpenGL graphics cores are available.
- Up to eight hours battery life with displays up to 8.9".
- Memory interface supports both DDR2 and mobile DDR1.
- MC13982 power-management integrated circuit, for reduced size and weight and more battery life.

The company is working first with Tier 1 OEMs (so watch for ASUS units) and expects products to start hitting the market in Q2 of this year—in other words, about now.

RESOURCES:

Freescale Semiconductor:

www.freescale.com/webapp/search/Serp.jsp

“New Freescale Processors Target Linux Netbooks”: arstechnica.com/journals/linux.ars/2009/01/07/new-freescale-processors-target-linux-netbooks

“Freescale Chip Aims at 1GHz, \$199 Netbook”: news.cnet.com/8301-1001_3-10130690-92.html

“Freescale, Nvidia Chase Netbook Sockets”: www.eetimes.com/news/latest/showArticle.jhtml?articleID=212700361

—DOC SEARLS



PHILIP JACOB

What They're Using:

PHILIP JACOB

I first met Phil Jacob on a consulting gig (yes, I got paid, but that was a long time ago), and quickly discovered that he was one of our most careful readers. Subsequently, he also became one of our most constructive critics—in the best sense of the word.

Phil is a born connoisseur. He applies that nature as founder and CTO of StyleFeeder (www.stylefeeder.com), a startup with a “personal shopping engine” that combines machine learning techniques and social data to provide users with a richer shopping experience. Among the features under StyleFeeder’s hood is one that personalizes search results from a dataset of more than ten million products using a custom-built recommendation engine. It also has a mature data tier shared across 100 MySQL databases in order to handle its scalability requirements. Given that StyleFeeder is

only six people, they’ve done a lot with a little during the past two years. Phil also is a technology advisor to PRX, a nonprofit licensing and distribution platform for public radio content. He’s also married to a well-known food blogger, Beatrice Peltre, who runs La Tartine Gourmande (www.latartinegourmande.com), so at least he eats well.

When I asked him to share some of what he’s using, he



concentrated his report on a single application—music:

Without a doubt, the device that attracts the most attention from guests visiting my house is the Slim Devices Squeezebox sitting in the living room. They usually don't notice it at first, but there's always that "double-take" moment when I change the music using my laptop or my iPod Touch.

I store my music in the basement on a Dell machine running Ubuntu with a bunch of FireWire disks attached for media files. This server runs Slim Devices' GPL-licensed SqueezeCenter software, which is a daemon that is mostly written in Perl, surprisingly enough. I happen to know Perl quite well as I programmed with it on a full-time basis for many years, so I was actually quite reluctant to buy this device when I first saw it. What I desperately wanted to avoid was a situation that required work just to get my music playing. And imagine a dinner party going silent because of a segfault in your own code!

I did, however, end up making a few changes to the software, which was only possible due to the open nature of it (at one point, I had reconfigured it to use MPlayer internally rather than LAME because

MPlayer wasn't as abusive to the CPU). One thing I really like about it is that I have a plugin installed that keeps a log of everything that gets played, which means that I could actually tell you what I was playing at 8:14pm last Thursday. One of my forthcoming weekend projects is to syndicate this onto my personal Web site, probably with a programming language that I'm looking to learn (OCaml is near the top of my list these days).

I access the music on my basement server when I'm at work by port-forwarding HTTP connections on my router to either Apache or to the SqueezeCenter's built-in streaming capabilities, depending on what I want to do. My router runs dd-wrt, a Linux-based firmware that runs on various types of home routers. It's very stable and provides a bunch of cool features that will appeal to power users.

Asked for last words on the subject, Phil adds, "We run RHEL and CentOS at StyleFeeder, some on EC2 and some at Contegix in St. Louis. We run behind Apache and nginx on the Web tier. We own zero servers and run dd-wrt on our office LAN as well....Typing one-handed with a baby in the other...." Babies are a new hack for Phil.

—DOC SEARLS

Expert included.

Johnny is a team leader in our hardware production facility, and he's responsible for operating system installations and validation testing. He is one of the experts who makes certain that when you use the Silicon Mechanics online configuration tools to order your server, you receive exactly what you specified, correctly configured and ready to perform. Johnny sized up the new 1U server from Silicon Mechanics—the Rackform nServ A108—and he's excited about its energy efficiency and cost effectiveness.

The AMD Opteron™ 1000 Series processors available with this server are offered in low power consumption models, and the A108 boasts an 86% efficient power supply. With a starting configuration price below \$750, the Rackform nServ A108 makes an outstanding entry-level server.

When you partner with Silicon Mechanics, you get more than an energy-efficient and cost-effective server configured just the way you want it—you get an expert like Johnny.



For more information about the Rackform nServ A108 visit www.siliconmechanics.com/A108.

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



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



REUVEN M. LERNER

Writing jQuery Plugins

If you're using jQuery, you already know JavaScript, so now it's time to write your own jQuery plugins.

The past two months, this column has looked at the jQuery library for JavaScript programming. jQuery is one of several popular libraries (like Prototype, YUI and Dojo) that have sprouted up in the last few years, making it possible to use JavaScript in ways that make the Web more satisfying and responsive by incorporating desktop-like behavior.

Part of the reason for jQuery's popularity is the huge library of plugins available for it. There are plugins for almost any type of functionality you can imagine—from GUI widgets to navigational aids to textual transformations. Plugins make it possible to isolate and reuse certain behaviors, achieving a goal known in the Ruby world as DRY (don't repeat yourself).

As I showed last month, using a plugin is generally quite easy. Download the plugin; install any CSS and JavaScript files that come with it, and then incorporate the JavaScript file into one or more HTML pages on your site, using a standard `<script>` tag.

Finally, attach the plugin to one or more elements on the page, using jQuery's event-handling functions, typically inserted into `$(document).ready`.

If you use jQuery, and you find yourself repeating the same JavaScript patterns over and over, you might want to consider writing your own plugin. Whether you distribute that plugin to the rest of the jQuery community depends on a number of factors, but by making it a plugin, you make it possible for all of your applications to load and use the library in a similar way.

Hubellubo Wuborld

A jQuery plugin is a packaging mechanism for your JavaScript code. This means in order to create your plugin, you first must have some JavaScript that needs packaging.

So, as an example this month, I've decided to create a simple translator into Ubbi Dubbi. Ubbi Dubbi, as some of you may know, is a "secret" language for children that was popularized in the United States by the public TV show *Zoom* in the 1970s (when I watched it), and then again in the 1990s. The rules for Ubbi Dubbi are simple. Every vowel (a, e, i, o and u) is prefixed with the letters ub. So, hello becomes hubellubo. It's not very hard to teach yourself to speak Ubbi Dubbi, and it sounds hilarious. Give it a try!

In any event, let's begin by creating a basic JavaScript program, using jQuery, that turns text into Ubbi Dubbi when the mouse cursor hovers over it. Let's start with a simple HTML file called `ubbi.html` (Listing 1). As you can see, there is no JavaScript in this file. Rather, we will use the "unobtrusive" style that jQuery encourages, writing our JavaScript in a separate file (`ubbi.js`, Listing 2), which we then include by means of a `<script>` tag.

The HTML itself is not very surprising or exciting. We have two paragraphs of text, each of which has the class `ubbi` assigned to it. In the JavaScript file, we use the `.ubbi` selector to set handlers for the mouseover and mouseout events. This is where the magic really happens. When the mouse hovers over the specified paragraph, the text is transformed into Ubbi Dubbi. When the mouse moves away, the text returns to its original form.

The translation depends on our `ubbify` function,

Listing 1. `ubbi.html`

```
<!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>
    <script type="text/javascript" src="jquery.js"></script>
    <script type="text/javascript" src="ubbi.js"></script>
    <link rel="stylesheet" type="text/css"
        media="screen" href="ubbi.css" />

    <title>Ubbi Dubbi</title>
  </head>
  <body>
    <h1>Ubbi Dubbi</h1>
    <p class="ubbi">This is in Ubbi Dubbi.</p>
    <p class="ubbi">
      Today, we will learn how to make cereal.
      First, pour the cereal into a bowl.
      Then pour milk onto the cereal.
      Finally, eat the cereal with a spoon. Delicious!
    </p>
  </body>
</html>
```


There are plugins for almost any type of functionality you can imagine—from GUI widgets to navigational aids to textual transformations.

which is defined as follows:

```
function ubbify(text) {
    return text.replace(/[aeiou]/gi, 'ub$1');
}
```

The above JavaScript function takes a single textual argument. It replaces any vowel with the string ub, followed by the letter that was replaced. Admittedly, there's a bug here related to capitalized words that begin with a vowel. Fixing that is left as an exercise for the reader.

Our mouseover handler is defined as follows:

```
$(".ubbi").bind('mouseover',
    function() {
        var original_text = $(this).html();
        $(this).attr({originalText: original_text});
        $(this).html(ubbify(original_text));
    });
```

This works by using jQuery's bind function, which invokes a function when a particular event fires on an HTML element (or collection of elements). So in this particular case, we tell JavaScript that every HTML element with a class of ubbi should invoke our function when the mouse cursor hovers over it. The function itself grabs the original text, puts it into an attribute named originalText, and then replaces the original text with the ubbified text.

The mouseout handler is similar, doing roughly the reverse, but without the ubbification:

```
$(".ubbi").bind('mouseout',
    function() {
        $(this).html($(this).attr("originalText"));
        $(this).attr({originalText: ""});
    });
```

To add a bit of pizzazz and styling, we also have ubbi.css, which uses the .ubbi:hover pseudo-selector to colorize and italicize the text when the mouse is hovering over it (Listing 3).

The combination of the CSS and JavaScript is fun and a bit exciting. Normally, the text looks as you would expect. But, when you move your

Listing 2. ubbi.js

```
function ubbify(text) {
    return text.replace(/[aeiou]/gi, 'ub$1');
}

$(document).ready(function() {
    $(".ubbi").bind('mouseover',
        function() {
            var original_text = $(this).html();
            $(this).attr({originalText: original_text});
            $(this).html(ubbify(original_text));
        });

    $(".ubbi").bind('mouseout',
        function() {
            $(this).html($(this).attr("originalText"));
            $(this).attr({originalText: ""});
        });
});
```

Listing 3. ubbi.css

```
.ubbi:hover {
    font-style: italic;
    border: 0.5px dashed #000;
    background-color: #cc9999;
}
```

mouse over a piece of text, it is transformed into Ubbi Dubbi. Prubetty cubo-ubol, rubight?

Making a Plugin

This JavaScript works just fine. However, perhaps there is a general need for Ubbi Dubbi translators that are active when the mouse hovers over text. It would be nice if someone simply could make every paragraph in a document automatically Ubbified with:

```
$(document).ready(function() {
    $(".p").ubbify();
});
```

In order to do this, let's create a jQuery plugin. The plugin, when incorporated, will add a new function to the jQuery object. This means that instead of our ubbify function being in the global namespace and instead of being invoked from within an event handler, we will define a function in the jQuery namespace, and it will

Normally, the text looks as you would expect. But, when you move your mouse over a piece of text, it is transformed into Ubbi Dubbi.

be invoked by handlers that also are defined in that namespace.

To make this happen, we need to restructure things a bit. First, we need to rename our JavaScript file, because every plugin needs to be of the format jquery.PLUGIN.js. In this case, I will call it jquery.ubbi.js.

Next, we need to define our ubbify function such that the global jQuery object will recognize it. To do this, we define ubbify inside the jQuery namespace:

```
$.fn.ubbify = function () {
    // implementation goes here
}
```

Listing 4. jquery.ubbi.js

```
(function($) {
    $.fn.ubbi = function(options) {
        // Private function
        function ubbify(text) {
            return text.replace(/([aeiou])/gi, 'ub$1');
        }

        // Return the results of iterating over our inputs
        return this.each(
            function() {
                $(this).bind(
                    'mouseover',
                    function() {
                        var original_text = $(this).html();
                        $(this).attr({originalText: original_text});
                        $(this).html(ubbify(original_text));
                    }
                );

                $(this).bind(
                    'mouseout',
                    function() {
                        $(this).html(
                            $(this).attr("originalText");
                        );
                        $(this).attr({originalText: ""});
                    }
                );
            }
        );
    };
})(jQuery);
```

Wait a second—what is this \$.fn that we are defining inside of? It turns out that if we want to define a global method for the jQuery object, normally aliased to \$, we must assign that function to the \$.fn object.

But, wait again—it is possible to redefine \$ so that it is no longer an alias to the \$ function. That allows jQuery to play nicely with JavaScript libraries such as Prototype, which also uses \$, but in a very different way. For this reason, many jQuery plugin tutorials tell you not to use \$, but rather the full jQuery object, like so:

```
jQuery.fn.ubbify = function () {
    // implementation goes here
}
```

Another solution is to wrap the entire function definition inside a closure (that is, a function with state), giving the closure the jQuery object as an environment with variable bindings:

```
($.fn.ubbify = function () {
    // implementation goes here
});
```

Now that we have gotten this out of the way, we can define our function inside its new plugin home. Listing 4 contains jquery.ubbi.js, a jQuery plugin that does everything we did before, but within the context of a plugin.

One of the most interesting things about jQuery is the fact that it accepts any number of arguments, thanks to CSS selectors. A function might be called for a single paragraph, identified via a DOM ID. Or, it might be invoked on many tags, or on tags with a certain class. Our function needs to handle any or all of these, and when it's done, our function must then return the jQuery object, so that its use can be “chained” to another set of instructions.

We do this by iterating over each argument and by returning the results, as follows:

```
return this.each(
    function() {
        ...
    });
```

jQuery defines .each to be an iterator that operates on each element of the object that invoked it. In this case, we take each of the submitted elements and pass them to a function. The function, of course, assigns the event handlers mouseover and mouseout. Notice how the functions are

One of the most interesting things about jQuery is the fact that it accepts any number of arguments, thanks to CSS selectors.

now invoked on `$(this)`, the jQuery version of the current element.

Finally, our `ubbify` function is defined privately within the `$.fn.ubbify` definition. Our `ubbify` function is available to any and all users within our definition of `$.fn.ubbify`, which is admittedly a very small number of functions for now.

With our plugin in place, all we have to do is tell our HTML file to load the plugin and to invoke it in the right way:

```
<script type="text/javascript" src="jquery.ubb.js"></script>
<script type="text/javascript">
    $(document).ready(function() {
        $(".ubb").ubb();
    });
</script>
```

Notice that `jquery.js` must be loaded before any plugins are loaded. We can apply our `ubbify` plugin to all of the paragraphs on a page with the following:

```
$(".p").ubb();
```

With our `Ubbi` plugin (`plubugubin?`) in place, it now has become that much easier to provide people with `Ubbi Dubbi` translations. Thanks to jQuery's plugin mechanism, we can distribute our plugin for others to use too, without having to read or understand the code. Our modified simple HTML file is shown in Listing 5.

Conclusion

jQuery is an amazing JavaScript library, but one of its particularly impressive features is support for plugins. Now that you have seen how easy it is to write a plugin, try to think of ways you can provide value to the community by publishing one or more plugins for others. ■

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.

Listing 5. ubbi2.html

```
<!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>
    <script type="text/javascript" src="jquery.js"></script>
    <script type="text/javascript" src="jquery.ubb.js"></script>

    <script type="text/javascript">
      $(document).ready(function() {
          $(".ubb").ubb();
      });
    </script>

    <link rel="stylesheet" type="text/css"
          media="screen" href="ubb.css" />

    <title>Ubbi Dubbi</title>
  </head>
  <body>
    <h1>Ubbi Dubbi</h1>
    <p class="ubb">This is not in Ubbi Dubbi.</p>
    <p class="ubb">
      Today, we will learn how to make cereal.
      First, pour the cereal into a bowl.
      Then pour milk onto the cereal.
      Finally, eat the cereal with a spoon. Delicious!
    </p>
  </body>
</html>
```

Resources

There are many resources on JavaScript and jQuery, both in print and on-line.

From Packt Press, I enjoyed *Learning jQuery* by Jonathan Chaffer and Karl Sweebber, which is good for Web developers who have experience in another language already, perhaps even JavaScript. It reviews many of the different types of functionality that a JavaScript programmer can accomplish using jQuery.

David Flanagan's *JavaScript: The Definitive Guide* continues to be an excellent resource, although I will admit that having jQuery has cut down significantly on what I need to know in the underlying JavaScript.

There are similarly many blog postings that might be helpful, including: www.learningjquery.com/2007/10/a-plugin-development-pattern, tkramar.blogspot.com/2008/02/improve-your-jquery-fu-write-plugins.html and www.bennadel.com/blog/800-My-First-jQuery-Plugin.htm.



MARCEL GAGNÉ

Backing Up to the Clouds

Enough with the cloud nonsense. Sure, you can buy in to some big commercial cloud, but odds are you've already got one of your own. Why not take advantage of it and get yourself backed up safely?

Our guests will be here any moment, François. What have you done with all the tables and chairs on the patio? Yes, I guessed you moved them to make room, which brings me to my second question. What are all these model rockets doing out here on the patio? Wait a minute. Are those USB Flash drives I see taped to the rocket bodies? Why, in the name of all that is free and open, are you planning to launch all those Flash drives? *Quoi?* What do you mean, you are doing what I asked you to do?

Mon Dieu! When I asked you to back up the restaurant's machines to our cloud, I didn't mean actual clouds. A *cloud* is just a fancy word for somebody else's network or network resources. More accurately, a cloud is a kind of marketing term to describe computing or storage resources provided by a vendor over the Internet. I was using the term in jest, partly because people seem rather obsessed with running their operating systems in the cloud, backing up their servers to the cloud and so on. The restaurant's network isn't on par with Google or Amazon, but we have plenty of storage space, *mon ami*. I see our guests arriving. *Vite, François!* Inside!

Good evening, everyone, and welcome to *Chez Marcel*, where exquisite wines are paired with the finest in Linux and open-source software. Please, take your seats and make yourselves comfortable while François cleans up his model rockets. We have a lovely 2006 Kim Crawford SP Flowers Sauvignon Blanc from New Zealand—a full-bodied dry white with intense flavor and a long finish...unforgettable. François, take those rockets to the cellar and return with the wine. *Merci, mon ami*.

While my faithful waiter fetches the wine, I probably should tell you that all those model rockets were an attempt at a cloud backup solution. There are, of course, numerous backup solutions you can use, from the simple and free to the complex and expensive, as well as everything in between. The technology behind most backup systems, however, tends to be much more limited. Using classic tools, such as tar and gzip, to back up and compress is still very common under the surface of much more

complex tools. This is true even when using network resources. In the end, you are backing up from one machine to another. Many people I know, including those with small businesses, do this for their regular backups. Machine A backs to machine B, which backs to C, which backs to A. The machines, and their drives, are all part of a network. Hey, instant cloud, and you probably didn't know you had one.

This is where rsync, another popular backup tool, shows its worth. As the name implies, rsyncs keep a backup copy of your data, in sync with the original. It can do it locally, from one physical drive to another, or across your network. Because only those files that have been modified are transferred, the process can be very quick. You can do this with single files, whole directories and subdirectories, while maintaining file ownership and permissions, links, symbolic links and so on. rsync has its own transport, or you can use OpenSSH to secure the transfer, and (of course) there are some great front-end, graphical tools to make the process a little slicker.

You can find rsync at rsync.samba.org, but you probably don't even have to look that far. Many distributions load it when you install your system. If not, check your installation disks or simply pick it up from your distribution's repositories. Before I explain how to rsync your data to your own personal cloud, let me show you how easy it is to create a synchronized backup of your data from one directory to another (or one drive to another):

```
rsync -av important_stuff/ is_backup
```

In the above example, rsync copies everything in the directory `important_stuff` into another directory (or folder) called `is_backup`. Most of you will have figured out that the `-v` means verbose copy. The `-a` option hides some amount of complexity in that it is the same as using the `-rlptgoD` flags. In order, this means that rsync should do a recursive copy; copy symbolic links; preserve permissions, modification times and group and owner information; and, with

the final D, copy special files (device and block). When you press Enter, files go scrolling by, after which you see something like this:

```

sending incremental file list
./
CookingJul08.tgz
CookingJul2008_albums.odt
CookingJul2008_albums.txt
igal_page.png
montage.png
shalbum.png
zenphoto_comment.png
zenphoto_go.png
zenphoto_login.png
zenphoto_makepass.png
zenphoto_setup.png
zenphoto_theming_comment.png
zenphoto_upload_photos.png
zenphoto_view_album.png
. . . .

sent 46059880 bytes  received 2753 bytes  6141684.40 bytes/sec
total size is 46044132  speedup is 1.00

```

One other thing that rsync should be able to do in order to be completely useful is delete files. If you are mirroring files and directories, it stands to reason that you want the mirror to represent exactly what is on the original. If files have been deleted, you want them deleted on the backup server as well. This is where the --delete parameter comes into play. Using the earlier example, let's delete that tgz file from the original, then relaunch the command:

```

$ rsync -av --delete important_stuff/ is_backup
sending incremental file list
./
deleting CookingJul08.tgz

sent 4164 bytes  received 25 bytes  8378.00 bytes/sec
total size is 41911050  speedup is 10005.03

```

From here on, both directories will always be in sync. When doing network backups, this magic synchronization of files and directories is done using a client and server setup. At least one machine must play the role of server (although nothing is stopping you from running an rsync daemon on every one of

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
- 80-320 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 4 GB RAM
- 80-320 GB hard drive / 128 GB SSD
- Pen/stylus input to screen
- Dynamic screen rotation
- Starts at \$2365

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 4 GB RAM
- 80-320 GB hard drive
- Call for quote

EmperorLinux

...where Linux & laptops converge

www.EmperorLinux.com

1-888-651-6686



Model specifications and availability may vary.

your machines). The server gets its information about who can access what from a configuration file called `rsyncd.conf`. You'll find that it probably lives in the `/etc` directory. The following partial listing is from one of my `rsync` servers:

```
hosts allow = 192.168.1.0/24
use chroot = no
max connections = 10
log file = /var/log/rsyncd.log
gid = nogroup
uid = nobody
```

```
[marcel]
  path = /media/bigdrive/backups/marcel
  read only = no
  comment = Marcel's files

[francois]
  path = /media/bigdrive/backups/francois
  read only = no
  comment = Files for the waiter
```

This configuration file is quite simple once you get the hang of it. Backup areas are identified by a name in square brackets (`marcel`, `website`, `francois` and so on). The chief bits of information there include the path to the disk area and some kind of comment. Notice that I specified `read only = no`, but I could just as easily have added that to the top section (the one without a name in square brackets). That's the global section. Anything put up there applies to all other sections, but it can be overridden. Pay particular attention to the `gid` and `uid` values; these are the group ID and user ID to which the file transfer takes place. The default is `nobody`, but you need to make sure that is correct for your system. One of my servers does not have a `nobody` group, but has a `nogroup` group instead.

The `hosts allow` section identifies my local subnet as being the only set of addresses from which transfers can take place. The `log file` line identifies a file to log information from the daemon. You also can specify a maximum number of connections, specific users who are allowed to transfer files (auth users) and a whole lot more. Run `man rsyncd.conf` for the full details. When your configuration is set, you can launch the `rsync` daemon, which, interestingly enough, is exactly the same program as the `rsync` command itself. Just do the following:

```
rsync --daemon
```

That's it. Now, it's time to put this setup to use. You might want to test your `rsync` connection by issuing the command:

```
rsync remote_host::
```

Note the double colon at the end of the server's name. The result should be something like this, assuming a server called `thevault`:

```
$ rsync thevault::
website    All our websites
francois   Files for the waiter
marcel     Backup area for Marcel
```

Now, pretend I am on the server where my Web site files live. Using the following command, I can launch `rsync` to back up this entire area:

```
rsync -av /var/www thevault::website/
building file list ...
```

The format of the `rsync` command is `rsync options source destination`, which means I also could start the command from `thevault`, assuming my Web site machine also was running an `rsync` daemon. The result would look more like this:

```
rsync -av localbackupdir websitemachine.dom::websites
```

All this work at the command line is great, but there are some tools for making the process easier, particularly if you will be creating a number of `rsync` backups or if you want to get into more complex requirements, such as scheduled backups. A friendly graphical front end on your desktop also may be a greater incentive to perform regular backups or take a

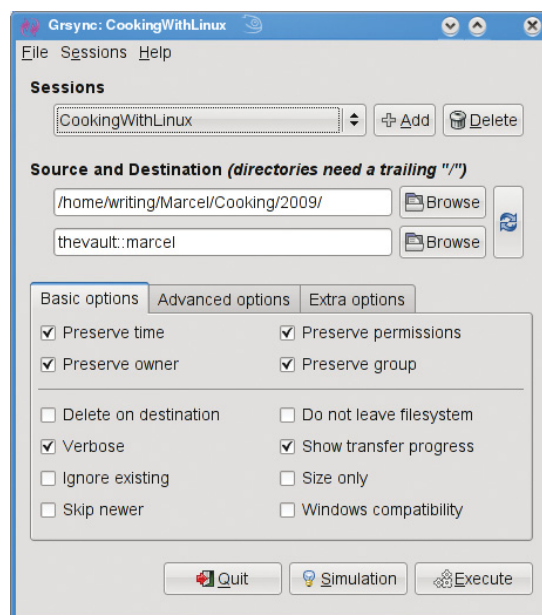


Figure 1. `grsync` provides an easy-to-use interface with every `rsync` option you could want.

quick backup when you've added important data and a "right now" backup is desirable. The first tool I want to show you is Piero Orsoni's grsync (Figure 1).

While providing a great front end to rsync, grsync also works as a teaching tool for the command-line version of the program, or at least it helps as a memory aid. Almost any command-line option available to rsync is covered in one of these three tabs: Basic options, Advanced options and Extra options. What makes it a learning tool is that if you pause over any of those check boxes with your mouse, a tooltip appears showing the command-line option with a brief description of its function.

To start, click the Add button next to the session drop-down dialog and enter a name for your backup. You can define many different rsync backups here, and then launch them again at a later time. Clicking the Browse button brings up the standard Gtk2 file browser window from which you can select your local and destination folders. Unfortunately, you can't browse remote systems, but if you've already set up an rsync server, have no fear. You can enter it manually in the format I showed you earlier (for example, `thevault::marcel/`).

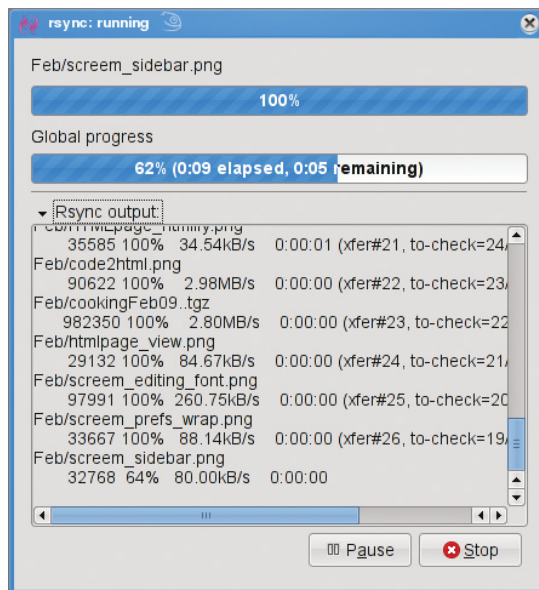


Figure 2. Once your grsync backup begins, it switches to a progress report view.

When you are happy with the various options, click Execute. If you only *think* you are happy, click the Simulation button. (Chef Marcel loves a

Expert included.

Jon is one of the most experienced sales experts on the Silicon Mechanics team, but he's noticed something new lately: Storform Storage by Silicon Mechanics is becoming **very** popular. Jon knows that his customers need to get the most for their money. They recognize real value in the storage servers and JBODs that he has to offer.

Storform Storage servers from Silicon Mechanics feature Intel® Xeon® Processor 5400 Series CPUs for fast, reliable compute power. They also offer build-to-fit options like 12 or 24 3.5-inch hot-swap drives in a 2U or 4U system, or 24 2.5-inch hot-swap drives in a 2U system. JBODs are also available, for uncomplicated scalability. With a starting configuration price below \$3650, it's no wonder Jon has noticed the rising popularity of these servers.

When you partner with Silicon Mechanics, you get more than great performance at affordable prices — you get an expert like Jon.



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

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

Pictured here are a few of our most popular Intel Xeon processor-based storage servers, from top to bottom: the Storform iServ R506, R505, and R503.

For more information about the Storform iServ line of storage servers, visit www.siliconmechanics.com/Storform.

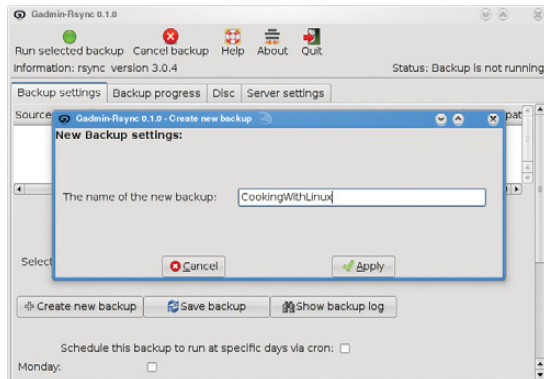


Figure 3. GAdmin-Rsync lets you define numerous backup configurations, each with its own identifier.

program with a sense of humor.) When you do click Execute, the program switches to a progress window (Figure 2), so you can see where you are in the process.

The next item on our rsync menu is Magnus Loeff's GAdmin-Rsync. GAdmin-Rsync makes every aspect of creating an rsync backup a matter of filling in the blanks. What's more, the program creates backups using SSH by default, which means you can set up rsync backups to any machine to which you have secure shell access. This also means you don't actually need to have an rsync daemon running on the remote machine if you have SSH access. Let me show you how it works.

When you start the program for the first time, you'll be asked for a name to give your new backup (Figure 3). You could back up the entire system or select specific folders of filesystems. Choose a name that makes sense to you based on what you want to back up. Enter a name, then click Apply to continue.

As you saw when we did this at the command line, rsync backups can be local, to a remote system or from a remote system. The next window looks




Figure 4. Your next step is to define the location of the backup.

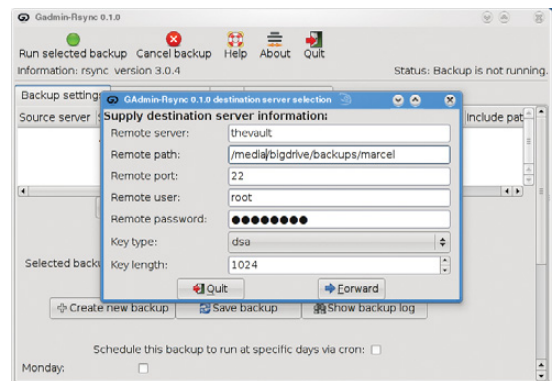


Figure 5. For remote backups, GAdmin-Rsync uses SSH/SCP for secure transfers.

for that very information (Figure 4). By default, local backup is checked. To back up to a remote server, select Local to remote backup. Because you can swap source and destination easily when using rsync, there's that third option. I routinely use a remote to local backup for my Web sites and remote systems. Click Forward to continue.

Assuming you chose to back up to your cloud, your next step is to enter the server information (Figure 5). This includes the backup path on your networked server as well as your SSH key type and length. When you have entered this information, click Forward.

Now you're ready to start the rsync backup. Click

TECH TIP Resetting the Root Password

The following methods can be used for resetting the root password if the root password is unknown.

If you use GRUB for booting, select the system to be booted, and add 1 to the end of the kernel boot command. If you're not presented with an edit "box" to add boot parameters, try using GRUB's edit command (the letter e). The 1 tells the kernel to boot to single-user mode.

The system now should boot to a root prompt. At this point, simply use the passwd command to change the root password.

Another option is to boot a rescue CD or an installation CD that lets you get to the command line. Once you're at

a command prompt, mount the system's root directory if it's not already mounted:

```
$ mkdir /mnt/system
$ mount /dev/sda1 /mnt/system
```

Now, do a chroot and reset the password:

```
$ chroot /mnt/system
$ passwd
```

—DASHAMIR HOXHA

Linux News and Headlines Delivered To You



Linux Journal
topical RSS feeds
NOW AVAILABLE

http://www.linuxjournal.com/rss_feeds

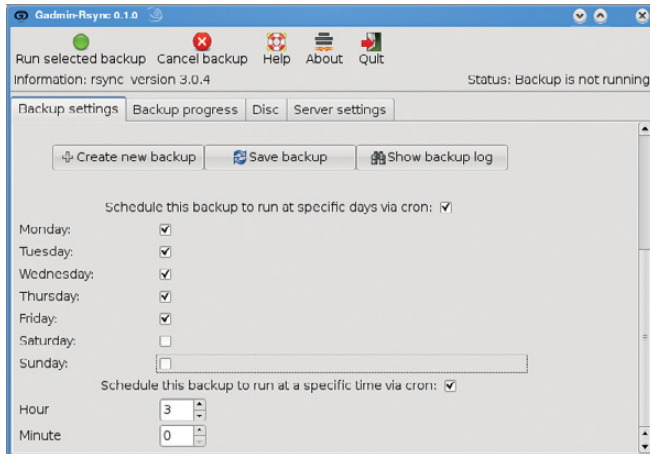


Figure 6. GAdmin-Rsync also provides an easy way to schedule your backups with cron.

the Backup Progress tab to watch all the action.

What is nice about this program is that you can (as with `grsync`) store a number of backup definitions, so you can choose to back up your documents, music or digital photographs when it suits you. GAdmin-Rsync goes further though. If you take a look down at the bottom of the window on the Backup settings tab, you'll notice the words "Schedule this backup to run at specific days via cron" and a check box (Figure 6). Check the box, then scroll down to choose the days you want the backup to run. A little further down, you can specify the time as well.

Well, *mes amis*, closing time has caught up to us, and at least for now, time is one thing we can't back up. Despite the hour, I am quite sure we can convince François to refill our glasses one final time before we go our separate ways. 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

GAdmin-Rsync: gadmintools.flippedweb.com

`grsync`: www.opbyte.it/grsync

`rsync`: rsync.samba.org

Marcel's Web Site: www.marcelgagne.com

Cooking with Linux: www.cookingwithlinux.com

WFTL Bytes!: wftlbytes.com



DAVE TAYLOR

More Fun with Word and Letter Counts

Determine the frequency of letters within a document (and become unbeatable at *Hangman*).

If you can remember back a month, you'll recall that I'd received a blessed e-mail from someone (hint, hint) asking:

Dear Dave, I seek an edge when I next play *Hangman* or other word games. I want to know what words are most common in the English language and what letters are most common in written material too. If you can show how to do that as a shell script, it'd be useful for your column, but if not, can you point me to an on-line resource? Thanks.—Mike R.

I grabbed three books from the Project Gutenberg archive ([gutenberg.org](http://www.gutenberg.org)) to analyze and use as test input: *Dracula* by Bram Stoker, *History of the United States* by Charles A. Beard and Mary Ritter Beard, and *Pride and Prejudice* by Jane Austen.

The obvious way to analyze these text files is with the `wc` command, which reveals that, combined, we're looking at 497,663 words, 2.9 million characters.

We used the following to identify the most common words:

```
$ cat *txt | tr ' ' '\012' | \
  tr '[:upper:]' '[:lower:]' | \
  tr -d '[:punct:]' | grep -v '^[a-z]' | \
  sort | uniq -c | sort -rn | head
```

The results were sufficient to reveal that the top ten words that appear in our 500,000-word sample are, in order: the, of, and, to, in, a, i, was, that and it.

Now, let's go in a different direction and analyze letter frequency. Then, we'll go back to finding interesting and unusual words.

Calculating Letter Frequency

The question underlying calculating letter frequency is this: "how do you break down a word into individual letters so that you have one letter per line?" It turns out that the handy Linux tool `fold` can do exactly what we want:

```
$ echo hello | fold -w1
h
e
```

```
l
l
o
```

Neatly done! (Note that you can't use `fmt` or similar commands because even if you specify `-w1` for width, it works with words, not characters.)

It's an easy leap from there to make `fold` break down every single word in a text file, sort the results, and use our power duo of `uniq -c | sort -rn` to get the results we seek:

```
$ fold -w1 < dracula.txt | sort | \
  uniq -c | sort -rn | head
157559
78409 e
56524 t
51608 a
50568 o
43453 n
41749 h
38150 s
37950 i
35001 r
```

A blank is the most common, but we can skip that visually rather than complicate our pipe with yet another process.

As I said in the beginning, E is the most common letter, but it's a surprise to see T as the second most common, frankly. Maybe it's because we're not compensating for upper-/lowercase? Let's try again:

```
$ fold -w1 < dracula.txt | sort | \
  tr '[:lower:]' '[:upper:]' | uniq -c | \
  sort -rn | head -5
157559
78409 E
56524 T
51608 A
50568 O
```

Wait a minute. We shouldn't get the same result! Hmm...can you see what I've done wrong? Hint: look at the order of commands in the pipe.

Got it? The `tr` needs to appear before the first

sort command, or it transforms the output, but after it already has been sorted separately. We also should strip out punctuation, which can be done with the tr command as well. Here's a better attempt:

```
$ fold -w1 < dracula.txt | \
  tr '[:lower:]' '[:upper:]' | sort | \
  tr -d '[:punct:]' | uniq -c | \
  sort -rn | head
157559
79011 E
58618 T
53146 A
51122 O
43975 N
43501 H
43423 I
39296 S
35607 R
```

Will this ordering change if we use all three of our books rather than just *Dracula*? Let's try it:

```
$ cat *.txt | fold -w1 | \
```

```
tr '[:lower:]' '[:upper:]' | sort | \
tr -d '[:punct:]' | uniq -c | \
sort -rn | head
```

```
468727
273409 E
201726 T
175637 A
169836 O
158561 N
155910 I
135513 S
133927 R
127716 H
```

Same result! In order of frequency, the letters appear in text in the following sequence: E T A O N I S R H D L C U M F W G P Y B V K X J Q Z. (I'm a bit surprised that J shows up so infrequently.)

You now know what order to guess letters in *Hangman*, if nothing else.

Speaking of Hangman

Before we wrap this up, let's go back through the words in our corpus and find just those that are at

MAYBE ONE DAY SECURITY WILL BE EASY. UNTIL THEN, THERE'S **BLACK HAT.**

Black Hat is a community that brings together security professionals and visionaries from everywhere: government, academia, corporate security and the underground. We exist to promote and improve information security by facilitating the timely sharing of accurate, actionable security knowledge.

Join us for a full slate of hands-on training sessions and four exciting tracks of presentations from the best minds in the information security space, all with a special focus on the needs of government and infrastructure security.

REGISTER NOW AT BLACKHAT.COM
USE CODE BHDC9LJMAG TO SAVE \$100 ON BRIEFINGS REGISTRATION



BLACK HAT DC 2009
FEBRUARY 16-19
ARLINGTON, VIRGINIA

Platinum Sponsor



nitrosecurity

Gold Sponsors

BlackBerry

NORMAN

GIGAMON

SAINT

IOActive

splunk

Looking Glass SYSTEMS

StillSecure

Microsoft

TRUSTED

NETWITNESS

least ten letters long and occur infrequently. Here's how I'll do that:

```
$ cat *.txt | tr ' '\012' | \
  tr '[:upper:]' '[:lower:]' | \
  tr -d '[:punct:]' | tr -d '[0-9]' | \
  sort | uniq -c | sort -n | \
  grep -E '.....' | head
1 abolitionists
1 accommodation
1 accommodations
1 accomplishing
1 accomplishments
1 accountability
1 achievements
1 acknowledging
1 acknowledgments
1 acquaintanceship
1 administrative
1 advertisement
```

That gives us long words that occur infrequently in the English language—or, at least only once in the 500,000-word corpus we've been analyzing. (True confession: I simply added more and more

dots to the grep regular expression until I weeded out almost all of the results. I could also have used `{10,}` to get ten-character or longer matches.)

Some of these words obviously are more common in everyday parlance than in these particular books, however, such as advertisement, which I'm sure occurs more than once every 500,000 words in normal conversation, or at least in the circles I frequent!

What would really be great for *Hangman* would be to apply the letter-frequency rule further, so that you extract the infrequently occurring words, then come up with a sum value for the frequency of each letter in the word (I'd assign E = 1, T = 2, A = 3, O = 4, for example) and identify the longest words with the highest scores. Those will be your very best *Hangman* words.

But, I'm out of space and last I checked, I was supposed to be writing about different variable reference formats in shell scripts anyway. I swear, next column, I'll get back to that. Unless you (hint, hint) write me a note with a puzzle or scripting challenge to solve. ■

Dave Taylor has been involved with UNIX since he first logged in to the ARPAnet in 1980. That means, 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.



7th Annual

2009 HIGH PERFORMANCE LINUX ON WALL STREET Show and Conference

April 6, 2009 (MONDAY)

Roosevelt Hotel, NYC

Madison Ave and 45th St, next to Grand Central Station

Low cost
Conference
\$295. Save \$100.
The Show is free.

Low Latency Data Centers featured.

High volatility and increased volume are creating more pressure on IT departments to seek out technology to lower budgets and increase agility and productivity.

Our Show is in the 2nd quarter of 2009 and many economists anticipate a turn-around in our US economy with President Obama and the 111th Congress mandating recovery legislation.

Wall Street and the financial markets will be looking for cost saving alternatives including High Performance, Linux, Open Source, Virtualization, Cloud Computing, Grid, Blade, among other technologies.

Our Show provides an efficient one-day demonstration showcase and networking opportunity for the New York financial markets in the heart of New York City.

2008 Sponsors



Visit: www.linuxonwallstreet.com

Keynote Speakers



Wim Coekaerts,
VP Linux Engineering,
Oracle



Inna Kuznetsova,
Dir. Linux Strategy,
IBM

Show Hours: Mon, Apr 6 8 - 4:00
Conference Hours: 8:30 - 4:50

Conference Producer: **A-TEAM**GROUP

Show Management:

Flagg Management Inc, 353 Lexington Ave
New York, NY 10016
(212) 286 0333 flaggmgt@msn.com

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**





MICK BAUER

Building a Secure Squid Web Proxy, Part I

Nurture your inner control freak with Squid.

Consider the venerable Web proxy—back when the Internet was new to most of us, setting up a Web proxy was a convenient way to grant users of an otherwise non-Internet-connected network access to the World Wide Web. The proxy also provided a convenient point to log outbound Web requests, to maintain whitelists of allowed sites or blacklists of forbidden sites and to enforce an extra layer of authentication in cases where some, but not all, of your users had Internet privileges.

Nowadays, of course, Internet access is ubiquitous. The eclipsing of proprietary LAN protocols by TCP/IP, combined with the technique of Network Address Translation (NAT), has made it easy to grant direct access from “internal” corporate and organizational networks to Internet sites. So the whole idea of a Web proxy is sort of obsolete, right?

Actually, no.

After last month’s editorial, we return to technical matters—specifically, to the venerable but assuredly not obsolete Web proxy. This month, I describe, in depth, the security benefits of proxying your outbound Web traffic, and some architectural and design considerations involved with doing so. In subsequent columns, I’ll show you how to build a secure Web proxy using Squid, the most popular open-source Web proxy package, plus a couple of adjunct programs that add key security functionality to Squid.

What Exactly Is a Web Proxy?

The last time I discussed proxies in this space was in my December 2002 article “Configuring and Using

an FTP Proxy”. (Where does the time go?) A quick definition, therefore, is in order.

The concept of a Web proxy is simple. Rather than allowing client systems to interact directly with Web servers, a Web proxy impersonates the server to the client, while simultaneously opening a *second* connection to the Web server on the client’s behalf and impersonating the client to that server. This is illustrated in Figure 1.

Because Web proxies have been so common for so long, all major Web browsers can be configured to communicate directly through Web proxies in a “proxy-aware” fashion. Alternatively, many Web proxies support “transparent” operation, in which Web clients are unaware of the proxy’s presence, but their traffic is diverted to the proxy via firewall rules or router policies.

Why Proxy?

Just because nowadays it’s easy to interconnect TCP/IP networks directly doesn’t mean you always *should*. If a nasty worm infects systems on your internal network, do you want to deal with the ramifications of the infection spreading outward, for example, to some critical business partner with whom your users communicate over the Internet?

In many organizations, network engineers take it for granted that all connected systems will use a “default route” that provides a path out to the Internet. In other organizations, however, it’s considered much less risky to direct all Web traffic out through a controlled Web proxy to which routes are internally published and to use *no default route whatsoever* at the LAN level.

This has the effect of allowing users to reach the Internet via the Web proxy—that is, to surf the Web—but not to use the Internet for non-Web applications, such as IRC, on-line gaming and so forth. It follows that what end users can’t do, neither can whatever malware that manages to infect their systems.

Obviously, this technique works only if you’ve got other types of gateways for the non-Web traffic you need to route outward, or if the only outbound Internet traffic you need to deal with is Web traffic. My point is, a Web proxy can be a very useful tool in controlling outbound Internet traffic.

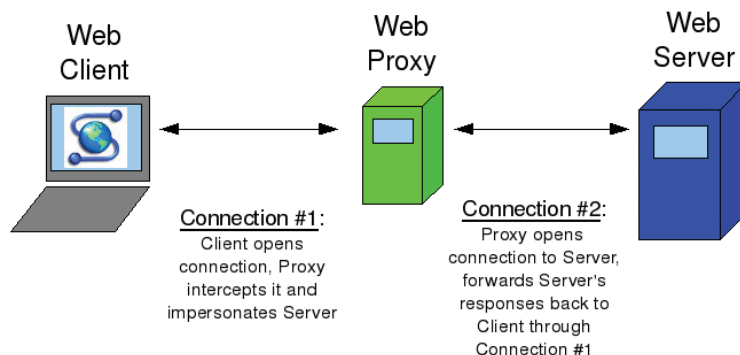


Figure 1. How Web Proxies Work

What if your organization is in a regulated industry, in which it's sometimes necessary to track some users' Web access? You can do that on your firewall, of course, but generally speaking, it's a bad idea to make a firewall log more than you have to for forensics purposes. This is because logging is I/O-intensive, and too much of it can impact negatively the firewall's ability to fulfill its primary function, assessing and dealing with network transactions. (Accordingly, it's common practice mainly to log "deny/reject" actions on firewalls and not to log "allowed" traffic except when troubleshooting.)

A Web proxy, therefore, provides a better place to capture and record logs of Web activity than on firewalls or network devices.

Another important security function of Web proxies is blacklisting. This is an unpleasant topic—if I didn't believe in personal choice and freedom, I wouldn't have been writing about open-source software since 2000—but the fact is that many organizations have legitimate, often critical, reasons for restricting their users' Web access.

A blacklist is a list of forbidden URLs and name domains. A good blacklist allows you to choose from different categories of URLs to block, such as social networking, sports, pornography, known spyware-propagators and so on. Note that not all blacklist categories necessarily involve restricting personal freedom per se; some blacklists provide categories of "known evil" sites that, regardless of whatever content they're actually advertising, are known to try to infect users with spyware or adware, or otherwise attack unsuspecting visitors.

And, I think a lot of Web site visitors do tend to be unsuspecting. The classic malware vector is the e-mail attachment—an image or executable binary that you trick the recipient into double-clicking on. But, what if you could execute code on users' systems without having to trick them into doing anything but visit a Web page?

In the post-Web 2.0 world, Web pages nearly always contain some sort of executable code (Java, JavaScript, ActiveX, .NET, PHP and so on), and even if your victim is running the best antivirus software with the latest signatures, it won't examine any of that code, let alone identify evil behavior in it. So, sure enough, the "hostile Web site" has become the cutting edge in malware propagation and identity theft.

Phishing Web sites typically depend on DNS redirection (usually through cache poisoning), which involves redirecting a *legitimate* URL to an attacker's IP address rather than that site's real IP, so they're difficult to protect against in URL or domain blacklists. (At any rate, none of the free blacklists I've looked at include a phishing category.) Spyware, however, is a common blacklist category, and a

good blacklist contains thousands of sites known to propagate client-side code you almost certainly don't want executed on your users' systems.

Obviously, no URL blacklist ever can cover more than a tiny fraction of the actual number of hostile Web sites active at any given moment. The real solution to the problem of hostile Web sites is some combination of client/endpoint security controls, better Web browser and operating system design, and in advancing the antivirus software industry beyond its reliance on virus signatures (hashes of known evil files), which it's been stuck on for decades.

Nevertheless, at this very early stage in our awareness of and ability to mitigate this type of risk, blacklists add *some* measure of protection where presently there's very little else. So, regardless of whether you need to restrict user activity per se (blocking access to porn and so forth), a blacklist with a well-maintained spyware category may be all the justification you need to add blacklisting capabilities to your Web proxy. SquidGuard can be used to add blacklists to the Squid Web proxy.

If you're serious about blocking access to sites that are inappropriate for your users, blacklisting is

Just How Intelligent Is a Web Proxy?

You should be aware of two important limitations in Web proxies. First, Web proxies generally aren't very smart about detecting evil Web content. Pretty much anything in the payloads of RFC-compliant HTTP and HTTPS packets will be copied verbatim from client-proxy transactions to proxy-server transactions, and vice versa.

Blacklists can somewhat reduce the chance of your users visiting evil sites in the first place, and content filters can check for *inappropriate* content and perhaps for viruses. But, hostile-Web-content attacks, such as invisible iframes that tell an attacker's evil Web application which sites you've visited, typically will not be detected or blocked by Squid or other mainstream Web proxies.

Note that enforcing RFC compliance is nothing to sneeze at. It constitutes a type of input validation that could mitigate the risk of certain types of buffer-overflow (and other unexpected server response) attacks. But nonetheless, it's true that many, many types of server-side evil can be perpetrated well within the bounds of RFC-compliant HTTP messages.

Second, encrypted HTTPS (SSL or TLS) sessions aren't truly proxied. They're tunneled through the Web proxy. The contents of HTTPS sessions are, in practical terms, completely opaque to the Web proxy.

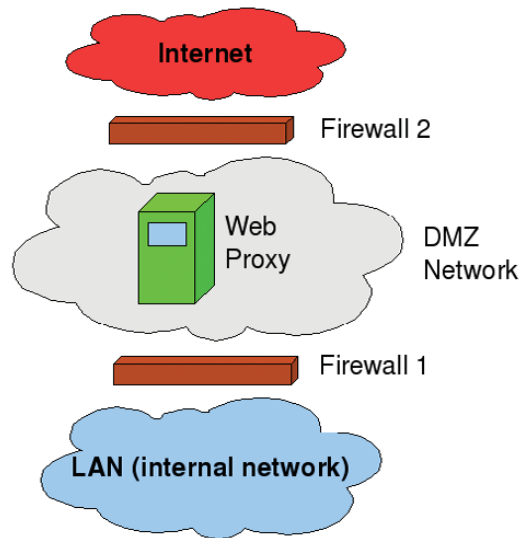


Figure 2. Web Proxy Architecture

an admittedly primitive approach. Therefore, in addition to blacklists, it makes sense to do some sort of content filtering as well—that is, automated inspection of actual Web content (in practice, mainly text) to determine its nature and manage it accordingly. DansGuardian is an open-source Web content filter that even has antivirus capabilities.

What if you need to limit use of your Web proxy, but for some reason, can't use a simple source-IP-address-based Access Control List (ACL)? One way to do this is by having your Web proxy authenticate users. Squid supports authentication via a number of methods, including LDAP, SMB and PAM. However, I'm probably not going to cover Web proxy authentication here any time soon—802.1x is a better way to authenticate users and devices at the network level.

Route-limiting, logging, blacklisting and authenticating are all security functions of Web proxies. I'd be remiss, however, not to mention the main reason many organizations deploy Web proxies, even though it isn't directly security-related—performance. By caching commonly accessed files and Web sites, a Web proxy can reduce an organization's Internet bandwidth usage significantly, while simultaneously speeding up end-users' sessions.

Fast and effective caching is, in fact, the primary design goal for Squid, which is why some of the features I've discussed here require add-on utilities for Squid (for example, blacklisting requires SquidGuard).

Web Proxy Architecture

Suppose you find all of this very convincing and want to use a Web proxy to enforce blacklists and conserve Internet bandwidth. Where in your

network topology should the proxy go?

Unlike a firewall, a Web proxy doesn't need to be, nor should it be, placed "in-line" as a choke point between your LAN and your Internet's uplink, although it is a good idea to place it in a DMZ network. If you have no default route, you can force all Web traffic to exit via the proxy by a combination of firewall rules, router ACLs and end-user Web browser configuration settings. Consider the network shown in Figure 2.

In Figure 2, Firewall 1 allows all outbound traffic to reach TCP port 3128 on the proxy in the DMZ. It does *not* allow any outbound traffic directly from the LAN to the Internet. It passes only packets explicitly addressed to the proxy. Firewall 2 allows all outbound traffic on TCP 80 and 443 from the proxy (and only from the proxy) to the entire Internet.

Because the proxy is connected to a switch or router in the DMZ, if some emergency occurs in which the proxy malfunctions but outbound Web traffic must still be passed, a simple firewall rule change can accommodate this. The proxy is only a logical control point, not a physical one.

Note also that this architecture could work with transparent proxying as well, if Firewall 1 is configured to redirect all outbound Web transactions to the Web proxy, and Firewall 2 is configured to redirect all inbound replies to Web transactions to the proxy.

You may be wondering, why does the Web proxy need to reside in a DMZ? Technically, it doesn't. You could put it on your LAN and have essentially identical rules on Firewalls 1 and 2 that allow outbound Web transactions only if they originate from the proxy.

But, what if some server-side attacker somehow manages to get at your Web proxy via some sort of "reverse-channel" attack that, for example, uses an unusually large HTTP response to execute a buffer-overflow attack against Squid? If the Web proxy is in a DMZ, the attacker will be able to attack systems on your LAN only through *additional* reverse-channel attacks that somehow exploit user-initiated outbound connections, because Firewall 1 allows no DMZ-originated, inbound transactions. It allows only LAN-originated, outbound transactions.

In contrast, if the Web proxy resides on your LAN, the attacker needs to get lucky with a reverse-channel attack only *once* and can scan for and execute more conventional attacks against your internal systems. For this reason, I think Web proxies are ideally situated in DMZ networks, although I acknowledge that the probability of a well-configured, well-patched Squid server being compromised via firewall-restricted Web transactions is probably low.

Yet to Come in This Series

I've explained (at a high level) how Web proxies work, described some of their security benefits and shown how they might fit into one's perimeter network architecture. What, exactly, will we be doing in subsequent articles?

First, we'll obtain and install Squid and create a basic configuration file. Next, we'll "harden" Squid so that only our intended users can proxy connections through it.

Once all that is working, we'll add SquidGuard for blacklisting, and DansGuardian for content filtering. I'll at least give pointers on using other add-on tools for Squid administration, log analysis and other useful functions.

Next month, therefore, we'll roll up our sleeves and plunge right in to the guts of Squid configuration and administration. Until then, be safe! ■

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

Resources

"Configuring and Using an FTP Proxy"
by Mick Bauer, *LJ*, December 2002:
www.linuxjournal.com/article/6333

The Squid home page, where you can obtain the latest source code and binaries for Squid:
www.squid-cache.org

The Squid User's Guide: www.deckle.co.za/squid-users-guide/Main_Page

The SquidGuard home page—SquidGuard allows you to enforce blacklists with Squid:
www.squidguard.org

The DansGuardian home page, a free content-filtering engine that can be used in conjunction with Squid: dansguardian.org

**METRO TORONTO CONVENTION CENTRE, CANADA
CONFERENCE & EXPO : 8 APRIL 2009**

**MAXIMIZE YOUR
BUSINESS SUCCESS**

it360°

UNIFIED COMMUNICATIONS
VoIP
SECURITY
DATA CENTERS

**Calling all IT professionals and business executives
to experience the latest leading industry advancements**

www.it360.ca

PRESENTING: **IN COOPERATION WITH:** **MEDIA PARTNER:** **ASTERISK & OPEN SOURCE TELEPHONY CONFERENCE SPONSOR** **MEDIA SPONSOR:**

ASTERISK AND OPEN SOURCE TELEPHONY CONFERENCE **CIO Executive Council**
Leaders Shaping the Future of Business **itWorld Canada**
IDG CANADA **digium** **LINUX JOURNAL**



KYLE RANKIN

When Disaster Strikes: Restoring a Master Boot Record

The Master Boot Record is an amazing little section of disk that you almost never notice—until it's gone. When that happens, read below to find out how to bring it back.

The following is a continuation of a series of columns on Linux disasters and how to recover from them, inspired in part by a Halloween *Linux Journal* Live episode titled “Horror Stories”. You can watch the original episode at www.linuxjournal.com/video/linux-journal-live-horror-stories.

I have to admit, I've learned more about how Linux works by breaking it and fixing it, than I have by any other method. There really is nothing quite like the prospect of losing valuable data, or the idea that your only computer won't boot, to motivate you to learn more about your system. In this month's installment of “When Disaster Strikes”, I discuss a surprisingly small part of your computer that plays a surprisingly large role in booting and using it—the Master Boot Record, or MBR for short. I cover some of my favorite ways to destroy an MBR and a few ways to restore it once you have.

Before you can fully understand how to restore the MBR, you should have a good idea of what it actually is. The MBR comprises the first 512 bytes of a hard drive. Now that's *bytes*, not megabytes or even kilobytes. In our terabyte age, it's hard to appreciate how very small that is, but to give you an idea, at this point in the column, I've already written about three MBRs worth of text.

This 512-byte space then is split up into two smaller sections. The first 446 bytes of the MBR contain the boot code—code like the first stage of GRUB that allows you to load an operating system. The final 66 bytes contain a 64-byte partition table and a 2-byte signature at the very end. That partition table is full of information about the primary and extended partitions on a disk, such as at which cylinder they start, at which cylinder they end, what type of partition they are and other useful data you typically don't think much about after a disk is set up—at least, until it's gone.

A Routine Lecture on Backups

This is the part of the column where I repeat some

of the best disaster recovery advice I know—make backups. In this case, we are talking about MBR disasters, so here are a few ways to back up your MBR. After all, it's only 512 bytes; there's no reason why you can't afford to back it up. Heck, it's small enough to tattoo on your arm, except I guarantee once you do you'll end up migrating to a new system or changing the partition layout.

The best tool to back up the MBR is coincidentally the best tool at destroying it (more on that later), `dd`. In fact, `dd` is one of those ancient, powerful and blunt UNIX tools that blindly does whatever you tell it to, and it's adept at destroying all sorts of valuable data (more precisely, it's adept at following your explicit orders to destroy your valuable data). The following command backs up the MBR on the `/dev/sda` disk to a file named `mbr_backup`:

```
$ sudo dd if=/dev/sda of=mbr_backup bs=512 count=1
```

Basically, this tells `dd` to read from `/dev/sda` 512 bytes at a time and output the result into `mbr_backup`, but to do only one 512-byte read. Now you can copy `mbr_backup` to another system or print it out and do the tattoo thing I mentioned before. Later on, if you were to wipe out your MBR, you could restore it (likely from some sort of rescue disk) with a slight twist on the above command. Simply swap the input and output sources:

```
$ sudo dd if=mbr_backup of=/dev/sda bs=512 count=1
```

More than One Way to Skin an MBR

There are a number of elaborate ways you can destroy some or all of your MBR. Please be careful with this first command. It actually deletes your MBR at the very least, and with a typo, it potentially could delete the entire disk, so step lightly. Let's start with the most blunt, `dd`:

```
$ sudo dd if=/dev/zero of=/dev/sda bs=512 count=1
```

This command basically blanks out your MBR by overwriting it with zeros. Now, unless you are masochistic, or you are like me and used this in a demonstration of MBR recovery tools, you probably wouldn't ever run this command. Most people end up destroying part of their MBR in one of two ways: mistakes with bootloaders and mistakes with fdisk or other partitioning tools.

Mistakes with partitioning tools probably are the most common way people break their MBRs, or more specifically, their partition tables. It could be that you ran fdisk on sda when you meant to run it on sdb. It could be that you just made a mistake when resizing a partition, and after a reboot, it wouldn't mount. The important thing to keep in mind is that when you use partitioning tools, they typically update only the partition table on the drive. Even if you resize a drive, unless you tell a partitioning tool to reformat the drive with a fresh filesystem, the actual data on the drive doesn't change. All that has changed are those 64 bytes at the beginning of the drive that say where the partitions begin and end. So, if you make a partitioning mistake, your data is fine. You just have to reconstruct that partition table.

It would figure that the first time I really destroyed my MBR, it was through the second, less-common way—mistakes with bootloaders. In my case, it was a number of years ago, and I was struggling to get an early version of GRUB installed on a disk. After the standard command-line commands didn't work, I had the bright idea that maybe I could use the GRUB boot floppy image. After all, it was 512 bytes and so was my MBR, right? Well, it sort of worked. GRUB did appear; however, what I didn't realize was that in addition to writing GRUB over the first 446 bytes of my MBR, I also wrote over the last 66 bytes, my partition table. So although GRUB worked, it didn't see any partitions on the drive.

Guessing Games Fix a Partition Table

I had at least used Linux long enough that after I made my mistake, I realized my actual data was still there and that there *must* be some way to restore the partition table. This was when I first came across the wonderful tool called gpart.

gpart is short for Guess Partition, and that is exactly what it does. When you run the gpart command, it scans through a disk looking for signs of partitions. If it finds what appears to be the beginning of a Windows FAT32 partition, for instance, it jots it down and continues until eventually it sees what appears to be the end. Once the tool has scanned the entire drive, it outputs its results to the screen for you to check and edit. It also optionally can write this reconstructed partition table back

to the disk.

gpart has been around for quite some time and is packaged by all of the major distributions, so you should be able to install it with your standard package manager. Don't confuse it with parted, which is a graphical partitioning tool. Of course, if your main system is the one with the problem, you need to find a rescue disk that has it. Knoppix and a number of other rescue-focused disks all include gpart out of the box.

To use gpart, run it with root privileges and give it the disk device to scan as an argument. Here's gpart's output from a scan of my laptop's drive:

```
greenfly@minimus:~$ sudo gpart /dev/sda

Begin scan...
Possible partition(Linux ext2), size(9773mb), offset(0mb)
Possible partition(Linux swap), size(980mb), offset(9773mb)
Possible partition(SGI XFS filesystem), size(20463mb), offset(10754mb)
End scan.

Checking partitions...
Partition(Linux ext2 filesystem): primary
Partition(Linux swap or Solaris/x86): primary
Partition(Linux ext2 filesystem): primary
Ok.
Guessed primary partition table:
Primary partition(1)
  type: 131(0x83)(Linux ext2 filesystem)
  size: 9773mb #s(20016920) s(63-20016982)
  chs: (0/1/1)-(1023/254/63)d (0/1/1)-(1245/254/56)r

Primary partition(2)
  type: 130(0x82)(Linux swap or Solaris/x86)
  size: 980mb #s(2008120) s(20016990-22025109)
  chs: (1023/254/63)-(1023/254/63)d (1246/0/1)-(1370/254/58)r

Primary partition(3)
  type: 131(0x83)(Linux ext2 filesystem)
  size: 20463mb #s(41909120) s(22025115-63934234)
  chs: (1023/254/63)-(1023/254/63)d (1371/0/1)-(3979/184/8)r

Primary partition(4)
  type: 000(0x00)(unused)
  size: 0mb #s(0) s(0-0)
  chs: (0/0/0)-(0/0/0)d (0/0/0)-(0/0/0)r
```

To hammer home the point about how easy it is to back up the MBR, now I have an extra backup of my laptop partition table—in this magazine.

As you can see, it correctly identified the two primary partitions (/ and /home) and the swap partition on my laptop and noted that the fourth primary partition was unused. Now, after reviewing this, if I decided that I wanted gpart to write its data to the drive, I would run:

```
$ sudo gpart -W /dev/sda /dev/sda
```

That isn't a typo; the `-W` argument tells `gpart` to which disk to write the partition table, but you still need to tell it which drive to scan. `gpart` potentially could scan one drive and write the partition table to another. Once you specify the `-W` option, `gpart` gives you some warnings to accept, but it also prompts you to edit the results from within `gpart` itself. Personally, I've always found it a bit more difficult to do it that way than it needs to be, so I skip the editor, have it write to the disk, and then use a tool like `fdisk` or `cdisk` to examine the drive afterward and make tweaks if necessary.

gpart Limitations

`gpart` is a great tool and has saved me a number of times, but it does have some limitations. For one, although `gpart` works very well with primary partitions, it is much more difficult for it to locate extended partitions, depending on which tool actually created them. Second, take `gpart` results with a grain of salt. It does its best to reconstruct drives, but you always should give its results a sanity check. For instance, I've seen where it has identified the end of a partition one or two megabytes short from the actual end. Typically, when we partition drives, we put one partition immediately after another, so these sorts of errors are pretty easy to find.

Reload the Boot Code

Now, if you have destroyed only the partition table, you hopefully should be restored at this point. If you managed to destroy the boot code as well, you need to restore it too. These days, most Linux

distributions use GRUB, so with your restored partition table, if you are currently booted into the affected system, run:

```
$ sudo grub-install --recheck /dev/sda
```

Replace `/dev/sda` with the path to your primary boot device. If you use an Ubuntu system, you optionally could use the `update-grub` tool instead. If you are currently booted in to a rescue disk, you first need to mount your root partition at, say, `/mnt/sda1`, and then use `chroot` to run `grub-install` within it:

```
$ sudo mkdir /mnt/sda1
$ sudo mount /dev/sda1 /mnt/sda1
$ sudo chroot /mnt/sda1 /usr/sbin/grub-install
  ↳--recheck /dev/sda
```

If the `chrooted grub-install` doesn't work, you typically can use your rescue disk's `grub-install` with the `--root-directory` option:

```
$ sudo /usr/sbin/grub-install --recheck
  ↳--root-directory /mnt/sda1 /dev/sda
```

Well hopefully, if you didn't have a profound respect for those 512 bytes at the beginning of your hard drive, you do now. The MBR is like many things in life that you don't miss until they are gone, but at least in this case, when it's gone, you might be able to bring it back. ■

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.

TECH TIP Stop Telling sudo Your Password

If you get tired of typing your password for `sudo`, but you don't want (or don't have permissions) to put `NOPASSWD` in your `sudoers` file, you can use the following procedure to update the `sudo` password timestamp and avoid typing your password.

Step 1) Create `$HOME/bin/sudo-hack.sh`:

```
#!/bin/bash
while [ true ];
do
    sudo -u root /bin/true > /dev/null 2> /dev/null
    sleep 60
done
```

Step 2) Do an initial run of `sudo` to set its password timestamp:

```
$ sudo -u root /bin/true
Password: *****
```

Step 3) Start `$HOME/bin/sudo-hack.sh` in the background:

```
$ HOME/bin/sudo-hack.sh &
```

Now you can use `sudo` without getting a password prompt, regardless of how long it's been since the last time you ran `sudo`.

Note: there are most certainly security implications related to using this procedure; of course, that's also true of using `NOPASSWD` in the `sudoers` file.

—PIETER DE RIJK

LinuxFest Northwest

Celebrating our 10th anniversary!



Grassroots Linux Gathering
Exhibits of all flavors
Presentations of all levels
Prizes, After parties
FREE Admission & Parking
FREE Open Source Software
Bring the whole family!

April 25-26 2009
Bellingham WA

Hosted By



Bellingham
TECHNICAL
COLLEGE

For schedules, event planning, and more, visit us at:

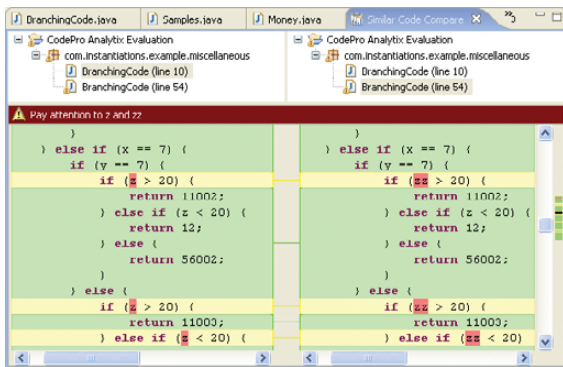

linuxfestnorthwest.org

Linuxfest Northwest is a Washington State Non-Profit Organization

SystemBase's Eddy Series Embedded CPU Modules

Hot and fresh from SystemBase's ovens are the new Eddy v2.1 Series embedded CPU modules for high-speed serial communication with real-time Linux. These small, Linux-ready 32-bit ARM9-based modules support high-speed RS-232, RS-422 and RS-485 serial interfaces at up to 921.6Kb/s, tolerate an extended temperature range (from -40°C to +85°C) and are equipped with Ethernet and wireless interfaces. SystemBase says that developers can set up their designs on the Eddy-DK v2.1 hardware development kit and the software development environment LemonIDE for Lemonix. Lemonix is an embedded, real-time Linux operating system that has been revised to support real-time capabilities while retaining the stable traits and merits of the Linux kernel 2.6.x.

www.sysbase.com



Instantiations' CodePro AnalytiX



If Java's your gig, take note of the latest v6.0 release of Instantiations' CodePro AnalytiX, a code-review tool for Eclipse Java. Instantiations asserts that CodePro AnalytiX's new product features will "help developers decrease potential code security vulnerabilities early in the software development life cycle, improve Java code quality and reduce development costs through increased developer productivity". Core product features include code audit, metrics, automated unit tests and team collaboration. New features include 25 new OWASP-based rules; two new audit-rule categories (Web services and threads and synchronization), mock objects to simulate the behavior of other objects safely and JUnit testing support for more Web application frameworks. JUnit test generation now supports most popular frameworks, including Spring, Struts and Enterprise Java Beans.

www.instantiations.com

Shai Vaingast's *Beginning Python Visualization* (Apress)

Although the title of Shai Vaingast's new book *Beginning Python Visualization* tells you something, the subtitle, *Crafting Visual Transformation Scripts*, perhaps tells you even more. Author Vaingast says that we are "visual animals" whose brains must sort, organize and transform data into images "before we can see the world in its true splendor". Part of Apress' Beginning Series, *Beginning Python Visualization* illustrates how to turn many types of small data sources into useful visual data. Learning Python is simply an added bonus. Readers will learn to set up and use an open-source environment as an alternative to Excel for data visualization. The book is for IT personnel, programmers, engineers, hobbyists and others who are interested in acquiring and displaying data from sources such as the Internet, sensors, economic trends, astronomical sources and more.

www.apress.com



Victor Kane's *Leveraging Drupal* (Wrox)

Thanks to past efforts of *LJ* founder Phil Hughes, Drupal holds a special place in the hearts of our editors. Thus, I will preach the good Drupal word by informing you of Victor Kane's new book *Leveraging Drupal: Getting Your Site Done Right*. Publisher Wrox calls *Leveraging Drupal* "much more than a tutorial" and a "nuts-and-bolts living mentor and guide" that explains how to do what is really required to build a site that works. Kane's book covers Drupal topics such as theming, customization and best practices in Web development while including videos and code and theme samples at every step. Advanced topics include views, panels and content creation.

www.wrox.com



ToutVirtual's VirtualIQ Pro

The company ToutVirtual bills itself elegantly as an “emerging leader” in virtualization based on its product VirtualIQ Pro. Now in Version 3, the product is “a single, platform-agnostic management and automation console” for virtualization deployment, says the company. VirtualIQ installs on Linux and Windows and provides features such as server-virtualization assessment, asset management, performance management, capacity management and reporting, together in one product. Users can support Xen, Citrix, Microsoft, Oracle and VMware virtualization platforms from one console. New features in version 3 include an updated UI, physical and virtual asset and inventory discovery, a physical-to-virtual migration analyzer and virtualization analytics.

A free version supporting up to five CPU sockets or 25 virtual machines is available for download from ToutVirtual's Web site.

www.toutvirtual.com

Blanco Kit

That Linux is the ideal OS for safeguarding data is logical. Meanwhile, the Blanco Kit from Blanco, Ltd., deploys Linux for exactly the opposite—to erase data safely from PCs. The recently upgraded Blanco Kit is a portable tool for performing data erasure at multiple sites, such as at dispersed warehouses storing obsolete computers. Blanco claims that “in a single day, just one engineer can erase 160 computers and share detailed reporting to meet compliance requirements”, as well as label the machines for resale or redeployment. The kit consists of a laptop PC with the Blanco Management Console and supporting applications, as well as a manual and label printer. The stylish aluminum case meets airline regulations as carry-on baggage.

www.blanco.com



Arkeia Network Backup

The new version 8.0 of Arkeia Network Backup is a backup and disaster recovery solution dedicated to supporting virtualized server environments. Version 8.0 extends Arkeia's deployment support for physical backup appliances and traditional software to support for virtual backup appliances, which the firm says is an industry first. The product also introduces backup and restore of complete virtual machine environments, permitting both traditional and virtual machine platforms to be protected with the same consistent process. The Arkeia Virtual Appliance delivers Arkeia Network Backup as a system image for a VMware virtual machine.

Other features in v8.0 include an improved graphical user interface, expanded reporting capabilities and integrated help system.

www.arkeia.com

Location Based Technologies' PocketFinder

From the “I wish I had thought of that” department comes Location Based Technologies' PocketFinder, a GPS-based personal locator system. The combination device and service allows users to check the real-time location of a PocketFinder device carried by a person, pet or suitcase at any time via the Internet, telephone or smartphone. The company calls its product “the smallest known single-board GSM/GPS device”. Advanced features include tracking of vehicle speeds and designation of customizable alert areas, whereby a notification is sent when a device enters or leaves a specified zone. The PocketFinder is also integrated into the iPhone and Google Android platforms.

www.pocketfinder.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

Ultrastar Deluxe—Home Karaoke Game

www.ultrastardeluxe.org

This project really caught my eye when browsing through SourceForge this month, and not being a big karaoke guy, this well-established project (more than 12 million users strong) has somehow passed me by all this time. For any others not in the know, *Ultrastar Deluxe* is “A free and open-source karaoke game inspired by the *Singstar* game available on the PlayStation. It allows up to six players to sing along with music using microphones in order to score points, depending on the pitch of the voice and the rhythm of singing.” *Ultrastar Deluxe* also runs on Linux, Windows and Mac OS X.




Ultrastar Deluxe brings you an amazing karaoke program with gorgeous graphics.

The most impressive feature I found was the ability to use pretty much any kind of video you want in the background, including MPEG, DivX, YouTube .flv—whatever you like. An extremely impressive and professional demonstration is available on YouTube (au.youtube.com/watch?v=uvBSeme34bE), where you can see the program in action with karaoke cues taking place on top of videos from artists such as the Foo Fighters and Tenacious D. *Ultrastar* isn't limited only to videos though. Standard gameplay generally takes place with an MP3 playing with a possible artist's picture in the background while the program takes information, such as timing cues and the like, from an accompanying text file. As the text file is raw ASCII and the format is well documented, it

leaves the entire process open to fans syncing up their own lyrical cues to any track they like, from any band. The game also comes with an editor built in, which will save hours of headache-inducing raw text editing.



Although I've got it set up with only one song, you can see how clean and colorful the interface is.



The guys from *Ultrastar* have put up a great YouTube video of their program in action. Here's someone rockin' out to the Foo Fighters.

Installation Installing *Ultrastar* can be a pain without someone holding your hand, but luckily there's a quick-start guide available that's well worth a look. The guide to installing it quickly under Ubuntu is very good, and it's probably the best guide for many other distributions too (if it doesn't work right away, or your distro doesn't use apt, at least the package names should help). I've chopped the Ubuntu guide down to the bare essentials that worked for me, but if you need more information, see the installation guide on-line (ultrastardeluxe.xtremeweb-hosting.net/wiki/doku.php?id=development:compile).

First, you need to satisfy *Ultrastar's* dependencies. Open a shell, and enter

the following:

```
$ sudo apt-get install fp-compiler fp-units-base
➤fp-units-misc fp-units-fcl fp-units-i386 fp-utils
➤libSDL1.2-dev libSDL-image1.2-dev portaudio19-dev
➤libsqlite3-dev libfreetype6-dev libavcodec-dev
➤libavformat-dev libswscale-dev
```

You can download the source either by SVN or wget, but for the sake of space, we'll just run with SVN here. If you don't have SVN, grab it with the following command:

```
$ sudo apt-get install subversion
```

Open a terminal where you want to save the *Ultrastar* source, and enter this command:

```
$ svn co
https://ultrastardx.svn.sourceforge.net/svnroot/
➤ultrastardx/trunk usdx-source
```

SVN will save the contents to the directory usdx-source. Enter that directory with this command:

```
$ cd usdx-source
```

And, compile it with these commands:

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

You can run *Ultrastar Deluxe* now, but I found with the SVN download, there weren't any songs added along with it (hence, nothing to load up and sing along to). If you check the project's Web site, there are some free songs to download provided by the community. Download a song's zip file and extract the contents somewhere locally. Then, with root or superuser permission, copy the files to the folder /usr/local/share/ultrastardx/songs. Once those are in place, you should be good to go.

Usage After all that, you now should be able to run the program with this command:

```
$ ultrastardx
```

Once you're in the game, you'll be

presented with the following menu options: sing, party, tools and quit. If you want to jump right into the action, choose sing. Enter your player name, select your difficulty level, and then you'll be taken to your song selection. Here you will find what are quite frankly gorgeous graphics, and an interface that should be familiar with users of programs such as iTunes and the like. Left and right arrows scroll through available tracks, and Enter starts playing a track. If you look at the bottom of the screen, there's a guide to further controls for better navigation and setting up a playlist. Once the track is playing, a number of cues will guide you quite easily through things such as timing and pitch, and anyone mildly familiar with karaoke should have no problems here.

Under party mode, you can choose from a number of innovative playing modes, such as the classic one on one with Duell. Blind Mode removes pitch cues leaving you to sing purely by feel. Until 5000 mode is where each player races to get 5,000 points, and the first one there wins. Once you've played for a while, it's worth checking what's available under tools, as there are options to tweak microphone settings, change the graphics resolution and change the in-game theme, among many others.

If you aren't getting anything from your microphone, check that it's working properly back in desktop land. It's probably a good idea to try to record a small amount of audio under a basic recording program and play it back to make sure everything's in order and sounding decent. Anyone who has the Singstar microphone from the PlayStation game should be able to use this too, according to the instructions (I don't have either the original game or mic, so I can't verify this).

Back in the game, there are settings to tweak the number of channels you can use, including the number of players, and even a party option with up to 12 players with three teams, but I don't have nearly enough space to cover that here.

Honestly, checking out the rest of the Web site is well worth it, as there's a bundle of themes available, including one that will make *Ultrastar* resemble the Nintendo Wii (quite appropriate given you also can use a

Wii controller). The links section of the Web site has instructions for a number of tools, such as the in-game editor and its syntax, plus there's a community board for requesting songs, exchanging help and so on.

Not only is *Ultrastar Deluxe* a great program for karaoke enthusiasts, it's also a great place for musicians to get their music out there to the public for free while having adoring teenagers singing along. Proprietary solutions may be better for getting famous songs on hand straightaway, but you would have to pay quite a premium for this, and the last time I tried karaoke, one of these machines was using a very lame midi soundtrack (which sounds awful when playing Metallica, trust me!)

If I were an enthusiast setting up at home and didn't mind scouting around for text files or spending some hours making your own, I'd use this program. If I were going into the karaoke business myself, I'd probably still use this program, given the amount of control you have and the ability to play original sound files. With the gorgeous graphics and well-thought-out interface, for karaoke, this is probably your best bet.

Editra—Advanced Text Editor

editra.org

For those looking for an advanced text editor that runs across the main three platforms of Linux, Windows and Mac OS, without being too bulky, this might be for you. According to the Web site: "Editra is a general, extensible, multiplatform text editor with an implementation that focuses on creating a clean and easy-to-use interface with features that aid in code development. Currently, it supports syntax highlighting and a variety of other useful features for more than 60 programming languages."

Installation I didn't run into any weird dependencies with this one, but it does require wxPython 2.8.3 or higher (which I already had installed). Although there are binary packages available only for Windows and Mac OS, there is a source package for Linux, and compiling Editra is pretty painless.

Head to the Web site, grab the latest tarball, and extract it. Open a terminal in the folder, and enter the following command either as root or sudo:

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

This built straight off for me, and it probably will give you little hassle, provided you have wx installed.

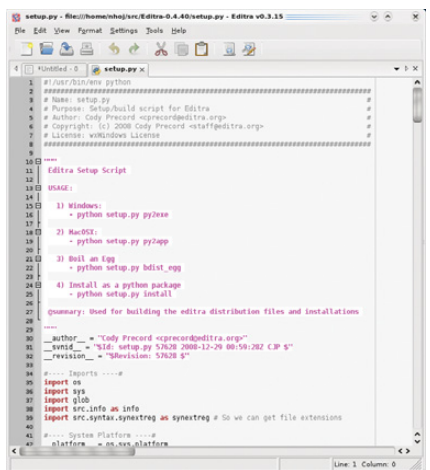
Usage Once it had compiled, I entered the command:

```
$ editra
```

That worked straightaway, and it even prompted me about an available update. It didn't work, but that's beside the point! That's some cool coding for something this early on in the development cycle. If you're reading along and have made it this far, you obviously don't need me to tell you how to use a text editor. However, let me draw attention to the cool features that sets Editra apart from other editors.

First, it's very lightweight and snappy. There's no bulkiness to put you off, and it probably would run nicely on old systems. Second, opening up extra files splits each one into tabs, like Opera, Konqueror, Firefox and so on. Third, it color-codes sections of text, depending on the language and syntax used. It recognizes a large number of languages too. I threw the following file types at it, and it knew Python, C, Pascal—heck, it even handled my configuration file for the X Window System.

It also has the extremely useful ability to open other character sets that are non-UTF-8, and it even suggests which encoding to choose to view the text you



```
1 #!/usr/bin/env python
2 #####
3 # Name: setup.py
4 # Purpose: Setup/build script for Editra
5 # Author: Cody Precond <precond@editra.org>
6 # Copyright: (c) 2008 Cody Precond <staff@editra.org>
7 # License: wxWindows License
8 #####
9
10
11
12 Editra Setup Script
13
14 USAGE:
15
16 1) Windows:
17    - python setup.py py2exe
18
19 2) Mac OS X:
20    - python setup.py py2app
21
22 3) Build an Egg:
23    - python setup.py bdist_egg
24
25 4) Install as a python package:
26    - python setup.py install
27
28 Summary: Used for building the editra distribution files and installations
29
30
31 author = "Cody Precond <precond@editra.org>"
32 version = "1.0: setup.py 5/10/2008 12:29 00:50:202 C/P S"
33 revision = "10Revision: 5/10/2008 S"
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
```

Editra brings a whole bunch of cool and sensible features into one snappy, multiplatform package.

Projects at a Glance

media hive—In-Browser Multimedia Player (www.mhive.org)



media hive

where within your browser. It works across multiple platforms, has a number of different interface modes, lets you rate the content of what you play and allows you to modify your playlist in your browser while it's streaming. It's also darn pretty to boot.

For people who want to enjoy music and videos on the go, media hive is a free program for playing multimedia from any-

Widelands—Real-Time Strategy Game (xoops.widelands.org)



Widelands

appears to have some pretty solid gameplay and impressive graphics for a free game. Hopefully, I can get it compiled and working and bring you an in-depth review in the future.

Although RTS games are generally not my favorite genre, *Widelands* draws its inspiration from such venerable classics as the *Settlers* series and

need (handy when jumping between systems and foreign countries). What really impressed me was the ability to “fold” paragraphs of text. Next to the line numbers (another handy feature), it groups text into specific bodies that are collapsible and expandable with that little box with a minus sign in it, which turns into a plus sign when collapsed. It's very handy for wrapping your head around lots of intricate code and keeping things organized, with the bonus that it also encourages clean coding.

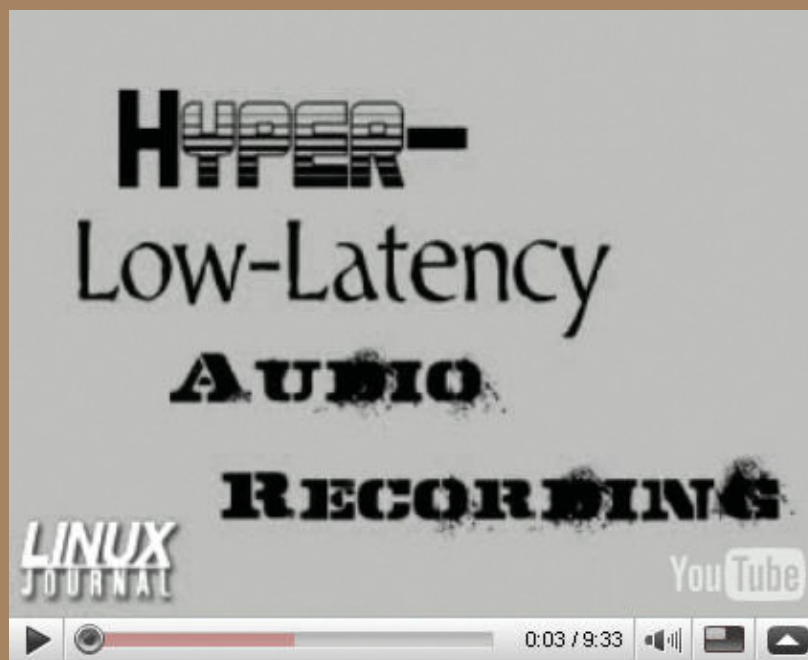
Editra has way more features available than what I've mentioned here, so check the Web site if you want to see the whole list. If you want an advanced text editor, you could do a lot worse than this. But, if you also want something that works across all of your systems—meaning you have to use only one interface—this might be the one. ■

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.

ON THE WEB, ARTICLES TALK!

The promise of zero or near-zero latency audio is a huge asset to the Linux operating system. Sometimes, achieving super low-latency audio is tricky, but not if your kernel is hard-real-time-capable. Cut your latency to less than 3ms with this tutorial:
www.linuxjournal.com/video/hyper-low-latency-audio-real-time-kernel



The Official Event of the Ruby on Rails Community



May 4–7, 2009

Las Vegas Hilton

www.railsconf.com

REGISTER NOW AND SAVE 15%
(Use discount code RC09LJ)



Co-presented by and Ruby Central, Inc.
O'Reilly Media, Inc.

O'REILLY
where2.0[™]
CONFERENCE

09 | MAY 19–21
SAN JOSE
CALIFORNIA


Becoming location aware

Now in its fifth year, the **Where 2.0 Conference** is where the grassroots and leading edge developers building location aware technology intersect with the businesses and entrepreneurs seeking out location apps, platforms, and hardware to gain a competitive edge.

Join with other developers, technologists, CTOs, researchers, geographers, academics, business developers, and entrepreneurs to debate and discuss what's viable now, and what's lurking just below the radar.

Where 2.0 — Your window into the geospatial industry and beyond

conferences.oreilly.com/where



REGISTER NOW & SAVE 15% USE DISCOUNT
CODE WHR09LNX

PXE

Not Just for Server Networks Anymore!

Using a combination of open-source technologies, you can build an unattended network-based OS installer that can save you huge amounts of time and even can install Windows.

BILL CHILDERS

In the April 2008 issue of *Linux Journal*, Kyle Rankin's article "PXE Magic" explains how PXE (Preboot eXecution Environment) works and how you can install your own PXE server and integrate rescue tools like Knoppix along with a PXE-capable Kickstart installation. I've used much of Kyle's PXE Magic before (he and I worked together in a previous life), but recently I found myself managing not only a network of Linux servers, but also the entire LAN, encompassing Ubuntu laptops, desktops and servers, along with Windows laptops,

desktops and servers. I found myself imagineering a PXE server that would not only allow me to kickstart servers and boot rescue tools off the network, but that also could provide a temporary environment for my users in the event that their computers broke. In my mind, the Holy Grail of this PXE server even would be able to install Windows machines via the network. After a fair amount of trial and error, I finally figured out the recipe, and in a strange twist, I was able to automate a network-based Windows installation...by bootstrapping Linux first.

Setting Up an Ubuntu Terminal Server

I knew one of my goals for this system would be to give the users of my network an environment they could PXE boot to in a pinch—something that would appear familiar to them, as well as allow them the ability to perform basic tasks like check e-mail, surf the Web, instant message and so on. Luckily, much of our staff here runs Ubuntu on the desktop, so the decision to implement an Ubuntu Terminal Server using the Linux Terminal Server Project (LTSP) was a simple one.

Like any PXE implementation, the LTSP server requires a TFTP server, a properly configured DHCP server and the syslinux software. In a nutshell, the client boots; the PXE code in the network adapter runs; the machine gets a DHCP address and the address of a server to grab the syslinux code via TFTP; and then, it actually runs a TFTP client and downloads that code and executes it, starting the boot process. Thanks to the hard work of the Ubuntu LTSP maintainers, setting up the server was fast and easy.

There are two paths you can take to install an LTSP server: normal or standalone. A normal LTSP installation assumes you have a pre-existing DHCP server on your network, and a standalone LTSP install assumes no DHCP server, and it will install the DHCP infrastructure and integrate it with the LTSP server automatically. There already was a DHCP server on our corporate LAN, so I elected to do the normal LTSP installation and integrate it with our existing Microsoft Windows DHCP server.

I began the installation by installing a standard Ubuntu 8.04 desktop on a Dell 1950 server, as the LTSP server will have to act as a GNOME desktop for anyone who would be logging in to it. After that, I assigned the server a static IP on our LAN (on the same subnet as the desktops and laptops). Installing the LTSP server was a piece of cake—a simple `sudo apt-get install ltsp-server openssh-server` at the GNOME terminal, and that task was complete. The final step on the LTSP server was to build the thin-client environment. Simply running `sudo ltsp-build-client` at the GNOME terminal fired off the remaining configuration steps and built the LTSP chroot.

Now that the LTSP server itself was ready, I had to

DHCP Notes

I mentioned integrating the LTSP server with a Microsoft Windows DHCP server, but it's not difficult to get the server to work with other DHCP servers. If you are running the "standard" ISC dhcpd server, see Kyle's "PXE Magic" article (listed in Resources). He includes example configuration options along with excellent explanations as to how they work. If you are running dnsmasq (popular in OpenWRT and other embedded or lightweight Linux distributions), the `dhcp-option=66,<ltsp_ip_address>` and `dhcp-option=pxe,67,pxelinux.0` in the `dnsmasq.conf` file should be what you need (I run this configuration at my home).

enable our network for PXE booting, and this meant messing with the Windows DHCP server. It took a little bit of trial and error, but much like in the DHCP server config that Kyle mentions in his article, there were only two configuration options that needed to be added to the DHCP scope. In Microsoft-ese, these were Option "066 Boot Server Host Name", which I set to the IP address assigned to the LTSP server and Option "067 Bootfile Name", which I set to "ltsp/i386/pxelinux.0". The last DHCP option seemed a little obscure, until I realized that the Ubuntu TFTP server's root directory was `/var/lib/tftpboot`. If you're running some other DHCP server, see the DHCP Notes sidebar, or refer to your DHCP server's documentation on adding options to the DHCP scope.



Figure 1. The Ubuntu LTSP GDM Login Screen

At this point, I could boot a PC on our LAN, press F12, select Onboard NIC as the boot device, and in about 30 seconds, I got a GDM login screen! I could log in to an LTSP session at this point, but I had to do it as one of the users that already was on the Ubuntu server. It was close, but not quite what I wanted, as the ideal setup would allow anyone on our Windows domain to log in to an LTSP session. Fixing this would have meant integrating the server with our corporate Active Directory. That used to be a major chore unto itself, but with Ubuntu 8.04 and higher, it's just an `apt-get` and a couple commands away.

The package that makes all this magic happen is called `likewise-open`. First, I ran:

```
sudo apt-get install likewise-open
```

to get the `likewise` package. After that, I had to get the Ubuntu server to "join" the Windows domain. I did this by running:

```
sudo domainjoin-cli join <fqdn.mydomain.com> <DomainAdminUID>
```

I wanted `likewise` to run when the machine boots, so I issued a:

FEATURE PXE

```
sudo update-rc.d likewise-open defaults
```

I also wanted the logins to be checked against the default domain, so I added the following line to the `/etc/samba/lwiauthd.conf` file:

```
winbind use default domain = yes
```

Finally, I started the likewise-open daemon using:

```
sudo /etc/init.d/likewise-open start
```

Now, my PXE LTSP clients could authenticate against the corporate Active Directory. Step one of the mission was complete!




Figure 2. The Ubuntu desktop—it's working!

Setting Up Unattended Windows Installs via PXE

The next step in my PXE adventures came when I was told I needed to refresh about 30 laptops with fresh builds of Windows. The method the previous Windows staff used to install Windows was through imaging the machine. Unfortunately, I could not locate the image files that had been used previously. Due to the issues and time constraints involved with trying to redevelop valid images for each hardware platform we had, I elected to do unattended installations of Windows.

I knew Windows included Remote Installation Service (RIS),

but because I was in a time crunch, I was reluctant to learn a completely new technology. However, there was another option: Unattended, an open-source project. I'd found the Unattended project about a year earlier, and although I'd dabbled with it in my home lab, I'd never tried it in a corporate environment. Like many Linux administrators, I hear "Windows" and I cringe, but because I was tasked with this, I figured I'd do my best to make sense of the Windows install process, as well as get some repeatability and understanding out of it.

Unattended relies on the fact that the first step of the Windows 2000/XP installer is essentially a DOS program. What happens when a machine is PXE booted to an Unattended install is a little convoluted, but it allows for great flexibility. Basically, the machine boots to a Linux kernel and shell, where some scripts provided by Unattended step in. The script partitions the system's disk and creates a basic FAT filesystem, and then it walks you through some menus where you can make choices as to the OS type (if you've set up Unattended with different Microsoft OS flavors), installation options and optional software you may have packaged. You're asked all the questions for a particular installation up front, including the CD Key, User Name, workgroup or domain to join, and administrative users. The Unattended scripts automatically digest all of this and create an `unattend.txt` file, which is dropped on the newly created FAT filesystem. Then, a FreeDOS session is started, and the Windows installer and OS bits are copied from a Samba share on the Unattended server, and then the installer is launched using the `unattend.txt` file. At this point, the installation is hands-off and proceeds without administrator intervention. The Unattended team has even gone so far as to create custom scripts that can install other pieces of software you may want to add to your configuration (automated VPN or Microsoft Office, for example).

Because there is no packaged install for Unattended, and the install process is quite different from the standard `./configure && make install`, I sat down for a bit and read the documentation on the site. Basically, the Unattended system leverages PXE and Linux as stated above, plus Samba for the distribution of the installation bits, and a bunch of Perl, shell and batch scripts to do a lot of the installation "magic".

The documentation asks that you have a working DHCP and DNS server, as well as a Samba server. I had the DHCP and DNS figured out for the LTSP server, so as per the step-by-step documentation, a `sudo apt-get install samba` got the Samba server installed. Next, I downloaded the Unattended distribution from the Web site and unpacked it in `/opt/unattended`. Then, I created a CNAME record on our DNS server that pointed `ntinstall` to the real hostname of the installation server. I then configured the Samba server with the following share information in `/etc/samba/smb.conf`:

```
[global]
...
guest account = guest
unix extensions = off
...
[install]
comment = Unattended
writable = no
```

WITH RESPECT TO AUTOMATIC UPDATE INSTALLATION, THE METHOD THE UNATTENDED FOLKS USE IS VERY LINUX-LIKE IN ITS RESOURCEFULNESS.

```
locking = no
path = /path/to/unattended/install
guest ok = yes
```

Finally, it was time to populate the OS distribution point with the Windows bits. This is done by creating a directory under the <unattended root>/install/os directory for whatever flavor of Windows you choose to install. In my case, I created an /opt/unattended/install/os/winxp directory and mounted that directory via Samba on my desktop. Then, I dropped the Windows installation media into the CD drive on my desktop and copied the /i386 directory from the CD to the /install/os/winxp share on the server. Now my Unattended server was, in theory, ready to install a system...except there was no way to select the Unattended install from a boot menu.

Here's where Kyle's article helped out again. What I needed was a PXE boot menu, and thanks to his article, I was able to whip up one in fairly short order. I had to take the bzimage and initrd files out of the tftpbboot directory in the linuxboot zip file on the Unattended site and place them in the /var/lib/tftpbboot/ltsp/i386/ directory (I renamed the bzimage to unat and the intird to unatin.img to help distinguish them better).

Then, I created a /var/lib/tftpbboot/ltsp/i386/pxelinux.cfg/default file (Listing 1) on the server, containing a combination of the syslinux boot arguments from the LTSP server and the Unattended server's configurations. Note the DISPLAY and LABEL directives. The DISPLAY directive states that when the machine boots you see the file pxemenu.msg displayed on the screen. This contains the text of the menu. The LABEL directive is what you type to boot a particular menu option. In this case, if I type "1", I get the Ubuntu LTSP session (this is also the default), and if I type "2", I get the Unattended Windows install.

Now, when I booted a PXE client, I got a choice as to whether to go to the Ubuntu LTSP session or the Unattended install. At this point, I tested the Unattended installation, and it sort of worked—it installed a base Windows system just fine, but it didn't install any of the drivers, nor any of the patches to the operating system. I realized just how spoiled I am by Ubuntu's driver coverage and update manager, but I slogged ahead and continued to work to refine the system so that the driver and

Listing 1. Example pxelinux.cfg/default file

```
default 1
serial 0,9600n8
timeout 300
prompt 1
DISPLAY pxemenu.msg
F1 pxemenu.msg

LABEL 1
KERNEL vmlinuz
APPEND ro initrd=initrd.img quiet splash

LABEL 2
KERNEL unat
APPEND initrd=unatin.img z_user=guest z_password=guest
z_path=//192.168.1.20/install
```



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



- 4 TB 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.



FEATURE PXE




Figure 3. Choosing What Software Goes on the Machine

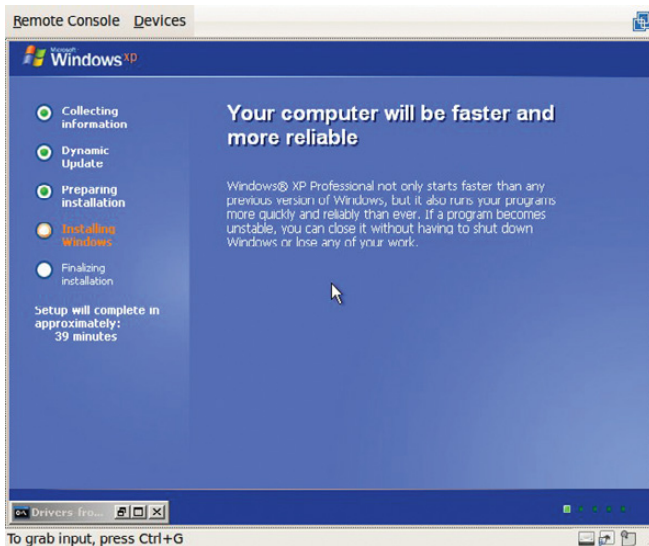


Figure 4. Windows installing! Faster and more reliable, that's debatable.

update installation happened without my intervention.

It turns out I didn't have to re-invent the wheel, as the driver issue and the update issue both have been addressed by the Unattended team. As far as the driver stuff goes, there is a method to integrate DriverPacks (which are large compressed archives of drivers) into the Unattended system. It's a little bit too involved for the scope of this article, but see the DriverPack link in the Resources section for more information.

With respect to automatic update installation, the method the Unattended folks use is very Linux-like in its resourcefulness. Under the Unattended root path, there are two directories: the `/install/scripts` and `/install/tools` directories. The scripts directory contains Windows batch files (.bat) that are used to do automated installation of various software packages, as well as some basic updates. The tools directory contains a set of scripts that will look at your Unattended server's current configuration and scripts directory, and then compare it to the CVS tree maintained by the Unattended team. The scripts then will grab the latest .bat files and drop them in the correct

place in the scripts directory. At that point, the next Windows install that's done with the Unattended system will get all the patches and install them automagically. The system even will reboot at the appropriate times, then pick up the next patch in the series and install it. To update the Unattended system's patch repository, it's as simple as running a `./script-update; ./check; ./prepare` from the `/install/tools` directory under the Unattended root.

The CVS archive of scripts, as well as the script archive on the wiki, proved to be invaluable. Those resources allowed me to finish the complete automation of my install, and now, I have a configuration that meets my company's needs for Windows. After about 30 seconds of typing the machine-specific information at the beginning of the installation, I now can walk away and know that Windows, Office, the Cisco VPN client, Symantec Anti-Virus and many other things my Windows users need will be done my way, automagically, without requiring myself or another staff member to babysit it.

In closing, thanks to the efforts of the Ubuntu and LTSP teams, I now have an environment that lets my users do some kind of work, even when their systems may have some kind of issue. And, thanks to the Unattended team, I don't have to sit at a Windows machine physically to install it, nor do I have to mess with half-baked images or other strange packaging solutions. I'm already getting other ideas on how to extend this system even further. ■

Bill Childers is an IT Manager in Silicon Valley, where he lives with his wife and two children. He enjoys Linux far too much, and probably should get more sun from time to time. In his spare time, he does work with the Gilroy Garlic Festival, but he does not smell like garlic.

Resources

"PXE Magic: Flexible Network Booting with Menus" by Kyle Rankin (April 2008 issue of *LJ*): www.linuxjournal.com/article/9963

Ubuntu Wiki—LTSP Installation: <https://help.ubuntu.com/community/UbuntuLTSP/LTSPQuickInstall>

Active Directory Authentication in Ubuntu 8.04 and 8.10: anothersysadmin.wordpress.com/2008/04/06/howto-active-directory-authentication-in-ubuntu-804

Unattended: a Windows Deployment System: unattended.sourceforge.net

Unattended Step-by-Step Instructions: unattended.sourceforge.net/step-by-step.php

Unattended Wiki: ubertechnique.com/unattended/FrontPage

Using DriverPacks with Unattended: ubertechnique.com/unattended/BTS_Driver_Packs

Unattended Script Archive: ubertechnique.com/unattended/Scripts

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

FreeBoo

an Open Architecture for Network Dual Boot

Using FreeBoo, you can restore and boot different operating systems across a network and replace proprietary solutions such as Rembo.

CRISTINA BARRADO AND SEBASTIAN GALIANO

Administrating large installations of computer desktops requires many tedious system reparations due to software updates, hardware fixes, user mistakes and viruses. To reduce costs, some enterprises adopt restrictive IT politics. But, if your business cannot afford a highly secure and restrictive environment, and you want to provide your many users with dual-boot capacity, desktop administration privileges

and the possibility to execute a large amount of different software, you probably are using Rembo.

This article presents FreeBoo, an open architecture that provides you with a dual-boot system of secure desktop images. FreeBoo is based on network boot, provides image restoration and allows hot boot. With FreeBoo, any malicious software installation done on a desktop by a previous user can be overwritten seamlessly.

Image Restoring and Dual Boot

Many IT departments' efforts are dedicated to the time-consuming task of repairing end-user desktops. For this task, most IT systems use open-source imaging systems that exist today, such as SystemImager, partimage, FileZilla, clonezilla, Frisbee, rsync, rdiff-backup, ADIOS and so on, or their commercial equivalents, including Norton Ghost, Active, True Image and Image. All these tools create a compressed image of a client's hard drive data and save it in case a future data recovery is necessary. An image is the complete copy of a filesystem, and it usually is stored on a backup server. When image changes are small, incremental backup is used to improve performance.

Imaging systems use well-known IETF protocols to transfer data from client to server or vice versa. They also include many functionalities for image management, with easy-to-use GUIs. IT departments in charge of large installations also use them to clone OS images onto several identical computers and to update systems with new patches.

In general, this software requires a high level of expertise, works basically on-demand and runs with a client program. This last feature is very important, because it assumes that the client computer is executing with specific conditions. Typically, this means the client always executes the same operating system.

At our university, computer labs can boot either MS Windows or Linux operating systems, and students select the desired partition using Rembo. Other PC-compatible dual-boot options include Norton BootMagic, OSL2000 or the MSTBOOT commercial systems and the GRUB open-source software solution. But, none of these tools can dual boot from the network.

Rembo is the only existing tool that provides the option of restoring any of the computer's saved images. And, very important, once the image is recovered, the computer boots it directly. Rembo is a commercial evolution of the open-source BP Batch Project, recently integrated into the IBM Tivoli suite. Rembo introduces local disk caches for fast image restoration, is able to use multicast messages and can be programmed using the Rembo-C scripting language.

FreeBoo Architecture

FreeBoo is proposed as an open-source, alternative solution to Rembo. Instead of open-source software, FreeBoo is an architecture built from many existing open-source programs. In fact, the number of new lines of code is insignificant. We have written only eight simple scripts and have used the urldecoder script authored by Heiner Steven. Figure 1 shows the global picture. The open protocols used include TFTP, DHCP, HTTP and NFS. And, the open-source code, running either on the server or client includes PXELinux, rsync, partimage, Apache, Netcat, links2 and gensplash. In Figure 1, you can see which program is executed at each computer site (client or server).

To illustrate the use of FreeBoo, let's assume the following scenario: ten or more desktops connected by LAN with a

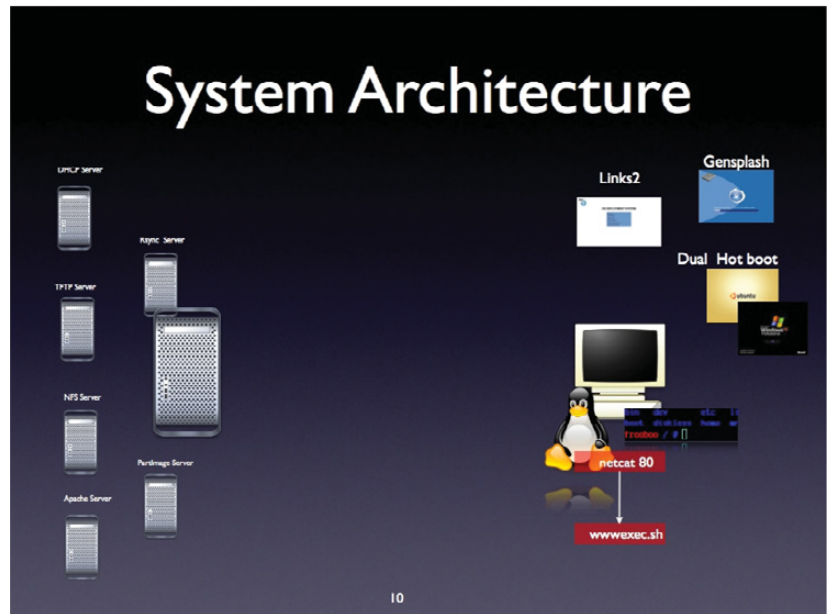


Figure 1. Global View of the FreeBoo Components

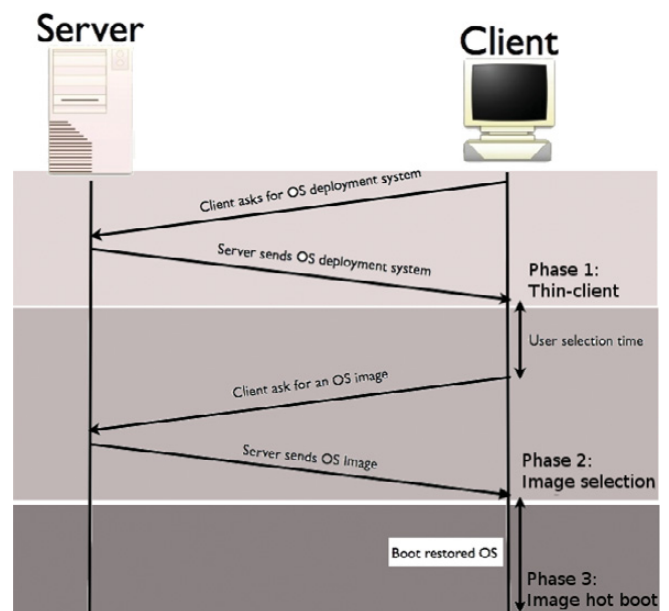


Figure 2. FreeBoo Phases

remote, non-accessible server room, using a PXE-compliant NIC. The desktops boot in three phases, as shown in Figure 2. The first phase is similar to a thin-client boot. In the second phase, the user selects a boot option, and the necessary image data is sent to the client. Finally, in the third phase, the client computer boots the user-selected OS.

Thin-Client Boot Phase

To build a FreeBoo system, you first need to configure your client's BIOS to boot from the network and to dedicate a server for PXE. You must start the DHCP and TFTP daemons on your server and use the PxeLinux.0 file as the PXE

primary bootstrap.

The main consideration regarding your thin-client configuration is the actual image properties. Because it is only a temporary image, let's configure a small embedded Gentoo Linux with generic drivers (less than 6Mb). We will add the FUSE drivers, the kexec system call introduced in the 2.6 kernel and NFS, rsync and partimage client utilities. FUSE is required, because we need write access to the NTFS partition during image restoration. kexec is used for the hot boot. The thin client uses an NFS root filesystem to avoid the initial RAM filesystem transfer done by most thin-client solutions. The NFS option is slower for individual access but faster for an initial deployment (more on using rsync and partimage in the next section).

To get a nice user interface, we also add the framebuffer driver, Gensplash (Gentoo's bootsplash software) and the Links2 browser compiled with direct framebuffer support. The Links2 text browser and Gensplash text images for framebuffer let you avoid the cost of the X Window System and its configuration problems, while achieving some graphical capacities.

Image Selection and Restoration Phase

The second phase consists of presenting a boot menu to the end user and (if required) executing the image restoration. This FreeBoo phase starts when the client PC is running the Linux thin client. You should configure it to execute the Links2 text Web browser initially. Simply add this `init.d` file to your thin-client filesystem (located in the server's drive and accessed via NFS):

```
#!/etc/init.d/freeboo
depends () {
    after gpm
    after local
    depends local
}
start() {
    /scripts/freeboostart.sh
    eend 0
}
```

The `freeboostart.sh` script executes the Web browser:

```
#!/bin/sh
# Part of the /scripts/freeboostart.sh file
links2 -g server_IP
# ...
```

The Apache Web server responds with an HTML file that contains the end-user menu. Figure 3 shows the available boot options. The user can choose to boot a given OS directly or to restore its saved copy. If the user selects a direct boot, FreeBoo skips directly to phase three. Otherwise, FreeBoo proceeds first to transfer the OS image to the client.

The menu provides two image restoring policies: complete (or full) restoring and partial (or fast) restoring. Complete restoring transfers all of the image to the client computer using `partimage`. Partial restoring transfers only the modified data using `rsync`. Depending on the differences with the

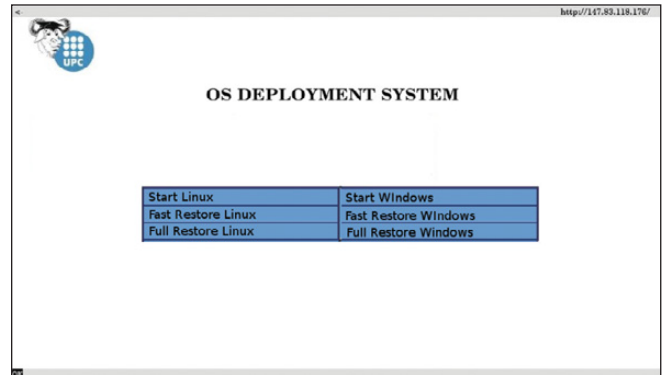


Figure 3. Boot Menu

current image, one solution is faster than the other. We have found that for partial restoring, the checksum computation was especially time consuming for large files. But, both solutions achieve the same goal—to deploy a clean and secure OS image on the client.

In addition to configuring Apache and setting the boot menu page as the user's initial page, you also need to start the `rsyncd` and the `partimaged` daemons on the server. `rsyncd`, for partial restoring, uses virtual names for the OS images. Its configuration file (`/etc/rsyncd.conf`) assigns each virtual name to a path in the filesystem. You should create a disk partition for each OS image, because `rsync` is filesystem-sensitive. Specifically, you need an NTFS filesystem and the `FUSE ntfs-3g` mount command for writing the MS Windows image on the client. We recommend excluding the big and unneeded MS Windows virtual memory file (`pagefile.sys`) and the hibernate file in the `rsync` recovery command (flag `--exclude-from`). We also recommend deleting any user-created files on the client (flag `--delete`). Additionally, add the flags `-a` (maintain permissions, ownership and timestamps), `-r` (recursive) and `-v` (verbose). Below is the code for the `syncroLinux.sh` and `syncroWindows.sh` scripts. They differ only in the mount command and image and partition names:

```
#!/bin/sh
# /script/syncroWindows.sh and /script/syncroLinux.sh
# FreeBoo scripts for partial image restoring
#
# INSERT the correct mount command
# for Linux:      mount /dev/hda2 /mnt/linux
# for MS-Windows: ntfs-3g /dev/hda1 /mnt/windows -o force

ntfs-3g /dev/hda1 /mnt/windows -o force

# INSERT the correct rsync command
# rsync -avr --delete --exclude-from=.rsync/exclude \
#     SERVER_IP::SYNC_NAME DEST_FOLDER

rsync -avr --delete --exclude-from=.rsync/exclude \
    192.168.1.1::windows /mnt/windows

# ... continues with kexec commands ...
```

Full Image Restoring

For full image restoration, we use partimage. partimage is much simpler than rsync, because it is not filesystem-dependent. You simply need to create a directory on the server to store the partimage chunk files. Chunks are just data that can be stored in any ext3 filesystem. Because the client gives the full path to the image chunks, you don't need any special configuration for partimaged. The flags to add are -f3 (quit when finished), -b (batch mode) and, for performance issues, -z# (level of compression). Below is the code that restores full images (the restoreLinux.sh and restoreWindows.sh scripts):

```
#!/bin/sh
# /scripts/restoreLinux.sh and /scripts/restoreWindows.sh
# FreeBoo scripts for full image restoring
#
# INSERT the correct partimage command
# partimage -f3 -b -s SERVER_IP restore DEST_DEVICE IMAGE
partimage -f3 -b -s 192.168.1.1 \
    restore /dev/hda2 /root/fileImage/linux

# ...continues with kexec commands ...
```

Hot Boot Phase

The third, and most challenging, FreeBoo phase is the image hot boot. When we started working on FreeBoo, our first idea was to enter into the Linux kernel code to write the hot boot function. But while examining the code, we discovered the kexec system

call and its related shell commands. The following two scripts are able to hot boot an OS image installed on a local drive.

Linux hot boot script:

```
#!/bin/sh
# /scripts/startLinux.sh
# FreeBoo script for Linux hot boot
#
mount /dev/hda2 /mnt/images/linux
kexec -l /mnt/images/linux/vmlinuz \
    --append="root=/dev/hda2 ro quiet splash" \
    --initrd="/mnt/images/linux/boot/initrd.img-2.6.15-23-386"
```

kexec -e

MS Windows hot boot script:

```
#!/bin/sh
# /scripts/startWindows.sh
# FreeBoo script for Windows hot boot
#
ntfs-3g /dev/hda1 /mnt/images/windows -o nonempty
kexec -l /mnt/images/windows/grub.exe
kexec -e
```

Both scripts first mount the partition of the OS image, then they execute kexec -l, and finally, they execute the new OS kernel with kexec -e. The OS kernel must be an ELF executable.

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.



www.logicsupply.com

To get a nice user interface, we also add the framebuffer driver, Gensplash (Gentoo's bootsplash software) and the Links2 browser compiled with direct framebuffer support.

For Linux, this is the kernel file directly, but for MS Windows, we use the GRUB bootloader (you should replace the NTLR default bootloader with the GRUB bootloader before saving the client MS Windows image on the server). Finally, you need a GRUB bootloader, such as Grub4dos, with built-in BIOS disk emulation and an ATAPI CD-ROM driver. This is needed because the thin-client execution overwrites the BIOS and most MS Windows versions still rely on some BIOS interrupts for video, timers and disk-related hardware I/O.

Script Execution

The final consideration to get FreeBoo working is in respect to the execution of the script files. We have one script file for each of the

six alternatives in the user's menu. These files are located on the server, but we need them to execute on the client. In the menu's HTML code, the links to the scripts are localhost references—for example, ``, so the Web browser will try to connect to a local Web server. But, we have not installed any Web server on the client image, because we wanted a light and fast kernel. Instead, we have included two additional scripts in FreeBoo to provide a local Web server. Both scripts are initiated by `/etc/init.d/freeboo`. Below are the complete `freeboostart.sh` and the new `mini_webapp_s.sh` scripts.

Initial `freeboostart.sh` script:

```
#!/bin/sh
# /scripts/freeboostart.sh
# FreeBoo script that starts a very simple http
# server with script execution capacity
# in parallel with a web browser
#
mini_webapp_s.sh &
links2 -g YOUR_SERVER_IP
```

The `mini_webapp_s.sh` script:

```
#!/bin/sh
# /scripts/mini_webapp_s.sh
# FreeBoo script that parses the incoming browser
# request, gets the script path name and
# executes it locally
#
torun=`nc -l -p 80 -s 127.0.0.1 | \
    awk '/HTTP/{print $2; exit}' | \
    urldecoder.sh`
($torun)
```

The first script starts the `links2` Web browser mentioned previously. But, it also starts the `mini_webapp_s.sh` script in parallel to act as a very simple Web server with application execution capacities. This second script executes the output of a pipeline command composed of `netcat`, `awk` and `urldecoder`, which extracts the filename of the script to execute.

`Netcat` (`nc`) is a very simple command. Like the traditional `cat` command, `netcat` simply copies data from an input stream to an output stream; the only difference is that these streams can be network data. The `-l` flag (listen mode) specifies that `netcat`'s input comes from the network. The `-p 80` and `-s 127.0.0.1` options indicate that the input will come on port 80 (the HTTP default port) from IP address 127.0.0.1 (localhost). `Netcat`'s function is to redirect any HTTP request, like the one below, to the `awk` filter:

```
GET /scripts/syncroWindows.sh HTTP/1.1
Host: localhost
User-Agent: ...etc...
```

The `awk` command extracts the script filename found on the HTTP GET line and passes it to the `urldecoder.sh` script through a second pipe. `urldecoder.sh` is a well-known script used to convert a URL with special characters, such as blank spaces, to a valid filename. The parser pipeline finishes when

Guests of Honor

Author Guest: Spider Robinson
-- Callahan's Crosstime Saloon

Media Guest: Wil Wheaton
-- Star Trek: TNG, Linux Geek, Author

Gaming Guest: Jane McConigal
-- Alternate Reality Games

Technical Guest: Rasmus Lerdorf
-- (creator of PHP)

Hack of Honor: Candyfab.org
-- prints 3D sculptures made of sugar

Nifty Guests

Jon "maddog" Hall -- Linux International

Steve Eley -- science fiction podcaster

Luke Ski & Tom Smith -- elite filkers

Eric S. Raymond -- OSS spokesperson

Also featuring, live singing tesla coil concerts and a steampunk masquerade.

Register today for a discounted rate!
<http://www.penguicon.com/>

May 1-3 2009 - Romulus, MI

PENGUICON 7.0

Science Fiction and Open Source Software Convention

the HTTP of the GET line is found. Then, the variable \$torun is set with its output and immediately executed. In the example above, the user has selected the MS Windows fast restore, and the HTTP request contains the /scripts/syncroWindows.sh filename. The pipeline extracts this name, and the next line executes it on the client.

The links2 browser, which is executing in parallel, is waiting for the HTTP response from our local Web server. Because we don't need any more interaction with the user, instead of sending back a response, we have included a line to kill the links2 process in each of the six menu scripts:

```
# first line of {start | syncro | restore}{Linux | Windows}.sh scripts  
killall links2
```

Conclusion

The BP Batch Project, started at the Geneve University by Marc V. Stuckelberg and David Clerc, became a popular open implementation of the thin client. This approach evolved into the commercial Rembo suite, which is used in many labs with a significant licensing cost. FreeBoo uses a combination of existing open-source technologies, including BP Batch, to provide the main features of Rembo.

The hardware requirements for installing FreeBoo are just a dedicated server connected to client desktops by a

LAN. Desktops need to have only boot-on-LAN capacities and local disk drives. All the software used is open source.

Future extensions of FreeBoo include the use of this technology for server software deployment; the development of a Web-based interface for easy administration of images, including database management; evaluation of the performance of the OS restoration process to improve it and to select the best option automatically (instead of having the user decide between the fast or full options); the insertion of multicast image recovery; and finally, the use of Wake-on-LAN capabilities to deploy secure images to desktops at preprogrammed times.

FreeBoo is only the initial step in building an open-source boot environment for system administrators that allows you to fix, deploy and execute OS images on large installations of desktops.

The scripts and other files related to FreeBoo can be found at ftp.linuxjournal.com/pub/lj/listings/issue180/10203.tgz. ■

Cristina Barrado is an Assistant Professor at the Technical University of Catalonia, in Castelldefels, Spain. She has been teaching operating systems since 1990 and has advised a large number of Master's theses on Linux and Linux development. Her PhD was focused on automatic extraction of low-level parallelism in loops at compile time. Currently, she belongs to the Icarus Research Group, whose target is research on avionics systems for Unmanned Aerial Systems.

Sebastian Galiano is a Telematic Engineer with extensive OS and network services knowledge. Currently, he is working at UPCnet as project engineer in the Internet and middleware area.



Calgary, Alberta
Canada

www.cossfest.ca

Calgary Open Source Systems Festival

Open Source for Business
Open Source for the Community

- Business Issues
- Workshops
- Wireless Security
- Programming
- System Administration
- Design
- Application Training
- Socialize!

April 17th & 18th
2009

Munin

the Raven Reports

Long-term monitoring with Munin is not restricted to system parameters. Why not monitor data of personal interest or data of interest to your colleagues?

PATRICIA JUNG

The long-term monitoring tool Munin was developed in Norway, and its name is derived from Norse mythology. One of the two ravens who report the news of the world to the god Odin is called Munin, and the other is named Hugin. Munin is “memory”, and Hugin is “thought”.

The cool thing about the Munin long-term monitoring suite is that it’s not restricted to supervising only typical system and network parameters. For example, you easily could do your colleagues from the marketing department a favor and monitor product sales for them. All you need to do is to write a script or program in the

language of your choice that returns the current value of the monitored parameter in the form `<parameter>.value <value>` (ending with a newline character) on the standard output.

It’s a little less trivial to understand how Munin actually works. This is because older parts of the official documentation (including the source code written in Perl) claim that Munin was a client-server application, which has caused much confusion. More recently, the Munin developers have referred to it as having a master-node architecture, which is far more appropriate.

A Master to Collect, Store and Present Data

To use the monitoring software, you need a machine that hosts a Web server and has Tobias Oetiker's RRDtool installed. On this computer, you run your Munin central, the Munin master. Debian and Ubuntu users will find it in the munin package. It consists of a set of Perl scripts to be run by the cron daemon in five-minute intervals. These scripts collect data from several Munin nodes, archive them in round-robin databases (RRDs), generate diagrams and update the Web pages that present them. One of them, dubbed munin-limits, also warns about values that violate limits, if configured accordingly.

RRDs have the advantage of never changing size. As time goes by, older data is squashed and stored in less detail until it finally phases out completely. Munin uses RRDs that store the data from the past 48 hours in five-minute resolution. Average values for the past ten days are stored with a resolution of half an hour; average values for the past 46 days are stored with a resolution of two hours, and average values for the past 449 days are stored with a resolution of one day.

The Munin master reads the munin.conf configuration file usually placed in /etc/ (under Debian/Ubuntu, /etc/munin/munin.conf) to find out where to ask for data and under which (host)name it should appear in the Web interface. A typical configuration entry looks like this:

```
[Airport;localhost.localdomain]
  address 127.0.0.1
  use_node_name yes
```

In brackets, you put the name of the machine to be monitored (here, localhost.localdomain). If you use a Fully Qualified Domain Name (FQDN), Munin automatically will present this machine as a member of a group named after the given domain part. If you prefer to use your own groups, add the relevant group name in front of the machine name in brackets (here, Airport; the result is shown in Figure 1).

Make sure you use a semicolon as the delimiter, without any whitespaces before and after, and decide on the group name before you start monitoring. This will make life easier for you, as Munin uses the group name as the name of the directory that contains the round-robin databases and the diagrams (in our example, on an Ubuntu system, /var/www/munin/Airport/).



Figure 1. Munin's main page presents all plugins ordered by group and, within a group, by hostname. Categories available for each host are presented in brackets.

The filenames of the databases and diagrams contain the hostname given in brackets. If you change the content of the brackets afterward, make sure to change file and directory names accordingly (and before the next five-minute interval is over); otherwise, Munin will use new empty RRDs, and you may lose data.

From Where to Obtain Data?

Use the address parameter to specify the Internet address from which the Munin master obtains the relevant data. This allows you to use an intermediate slave machine to gather data from the actual target machines without giving misleading information in the Munin Web interface.

What seems to complicate matters at first glance is actually a very useful feature, as it allows you to restrict the plain-text communication of the Munin protocol to trusted machines. In addition, you do not always have the opportunity to install the Munin node software on the actual target machine. This is, for example, the case if a Munin node gathers data via SNMP. In this case, you must set the use_node_name parameter to no.

On each machine given as the value for an address parameter, you need to install the Munin node software; otherwise, the Web pages generated by the Munin master will remain empty. On Debian/Ubuntu systems, the relevant package is called munin-node.

The Munin node consists of a daemon that, on request from the Munin master, starts the plugins responsible for collecting their specific type of data. Its configuration file, munin-node.conf, is stored in the same directory as the munin.conf; don't confuse the two if your Munin master runs on a machine that also acts as a node. munin-node.conf defines, among other things, the log file and log level, the port to be used (usually 4949), and most important, the machines that are allowed to connect to the daemon:

```
allow ^127\.0\.[0-9]\.1$
```

In this case, the regular expression ^127\.0\.[0-9]\.1\$ restricts access to the Munin master running on the same machine, localhost.

First Contact via Telnet

Installing the Munin node from a distribution package usually will activate a range of plugins that can be configured automatically. In this case, Telnetting to port 4949 of the node machine will give you an overview:

```
$ telnet localhost 4949
Trying 127.0.0.1...
Connected to localhost.localdomain.
Escape character is '^]'.
$ munin node at extrablatt.trish.de
help
$ Unknown command. Try list, nodes, config, fetch, version or quit
list
  open_inodes      if_err_eth0      irqstats
  entropy          processes        postfix_mailqueue
  if_eth0 df         netstat          interrupts
  swap             load             cpu
  df_inode         if_eth1          if_err_eth1
  postfix_mailvolume forks             iostat
  open_files       memory           vmstat

fetch open_inodes
used.value 67839
max.value 68094
.
Connection closed by foreign host.
```

The list command returns the names of all activated plugins on this particular node. The fetch command, with the name of the plugin you want to run as an argument, returns the values of the parameters (in this case, used and max) the plugin monitors. As long as you don't change the timeout directive in munin-node.conf (for example, to 20 seconds: timeout 20), you have to be fast to type in your commands, as the daemon will close the connection after a default of ten seconds.

If the list command does not list any plugins (which might be the case after a source code installation), you need to activate them first. To do this, symlink them into the plugins directory (/etc/munin/plugins on Debian/Ubuntu) on the node machine and restart the daemon using its init script. The command munin-node-configure --shell will show you link commands for some of the plugins provided with the distribution.

Wild-Card and autoconf Plugins

There are two types of plugins. For plugins that are independent of additional parameters, the link name equals the plugin name. Sometimes, however, a plugin can monitor several items of the same type, for example, several network interfaces, such as eth0 and eth1.

In this case, it would be stupid to hard-code the interface name into the plugin. Instead, you provide this information in the name of the symlink. Plugins capable of this are called wild-card plugins, and their names end with an underscore. If you, for example, want to monitor the eth0 interface with the wild-card plugin if_, the link pointing to the if_ plugin would be if_eth0:

```
$ ls -al /etc/munin/plugins/if_eth*
lrwxrwxrwx 1 root root 28 2008-06-27 23:53
->/etc/munin/plugins/if_eth0 -> /usr/share/munin/plugins/if_
lrwxrwxrwx 1 root root 28 2008-06-27 23:53
->/etc/munin/plugins/if_eth1 -> /usr/share/munin/plugins/if_
```

Almost all plugins provided with the Munin distribution belong to the plugin family auto and can be run with the autoconf argument. In this case, they check whether they are able to provide meaningful results. For example, running a monitoring plugin for the Exim MTA makes sense only if you're running the Exim mail server. If your system does not fulfill the prerequisites to run a specific autoconf-enabled plugin, it will provide you with meaningful hints:

```
$ /usr/share/munin/plugins/exim_mailqueue autoconf
no (exim not found)
```

To get an overview of all preinstalled plugins that implement the autoconf method, simply run:

```
# munin-node-configure --suggest
Plugin          | Used | Suggestions
-----
[...]
exim_mailqueue  | no   | [exim not found]
[...]
if_              | yes  |
[...]
```

If you want to share your own plugins with others, for example, at the Munin Exchange platform, we recommend you have a look at the officially provided plugins to see how the autoconf method is implemented. But, to get started with your own plugins, don't complicate your life unnecessarily.

Fly Away

As an example, let's monitor the departures at Munich Airport in five-minute intervals, as that is the default for the Munin cron job. The departure timetable is available from a Web page, and we can use a shell script and the Links text browser to dump it into a temporary file:

```
#!/bin/sh

SITE=http://www.munich-airport.de
DEP_URL=$SITE/en/consumer/fluginfo/abflug/index.jsp?viewType=t
TMP_FILE=/tmp/.muc_flights

links -dump $DEP_URL > $TMP_FILE
```

This file now contains lines like these:

```
[...]
[ LH 3464 ] [ Budapest ] [ 21:30 ] [ ] [ T2 ] [ departed ]
[ LH 726 ] [ Shanghai ] [ 21:30 ] [ ] [ T2 ] [ boarding ]
[...]
```

The Web page lists five flight states: calling, boarding, departed, planned (which means delayed) and cancelled. We will count them for the current time interval and return them on the standard output, like this:

```
calling.value 0
boarding.value 1
departed.value 1
planned.value 0
cancelled.value 0
```

To determine the current time, we could use the following date command:

```
$ date +%H:%M
21:30
```

and grep for this string in the "links" dump. Unfortunately, matters are a bit more complicated. The airport's departure timetable lists all flights in five-minute slots. But, even though the cron job on the Munin master is configured to run at zero, five, ten (and so on) minutes past the hour, we can't be sure it will run exactly on time. That's why our plugin uses a modulo operation (executed by bc) to round down the current minutes accordingly and combine hour and minutes in a case construction:

```
TIME=$(date +%H)
MIN=$(echo "($(date +%M)/5)*5"|bc)
case $MIN in
0) TIME=$TIME:00 ;;
5) TIME=$TIME:05 ;;
*) TIME=$TIME:$MIN ;;
```

esac

Now the TIME variable contains only hour:minute pairs in which the minutes are a multiple of five.

Stop Interpolating!

There's another complication—when the Munin master stores the values retrieved from the plugin in the relevant RRDs at a point in time deviating from the exact planned time that RRDtool interpolates them. This way the database rarely contains the integer values our plugin returns but slightly deviating floating-point values.

The good news is that RRDtool accepts time value pairs. In this case, it will refrain from interpolation. The time must be given in seconds since January 1, 1970 00:00:00 UTC as a prefix of the value. A colon (:) is used as the delimiter, as in the following example plugin output:

```
calling.value 1230841800:0
boarding.value 1230841800:1
departed.value 1230841800:1
planned.value 1230841800:0
cancelled.value 1230841800:0
```

(1230841800 equals January 1, 2009, 21:30.) Note that Munin versions before 1.3.4 were unable to handle plugin output using this extended format. This means the following plugin code won't be compatible with older Munin versions:

```
links -dump $DEP_URL | grep $TIME > $TMP_FILE
```

```
UNIXTIME=$(date -d$TIME:00 +%s)
echo "calling.value $UNIXTIME:$(grep calling $TMP_FILE | wc -l)"
echo "boarding.value $UNIXTIME:$(grep boarding $TMP_FILE | wc -l)"
echo "departed.value $UNIXTIME:$(grep departed $TMP_FILE | wc -l)"
echo "planned.value $UNIXTIME:$(grep planned $TMP_FILE | wc -l)"
echo "cancelled.value $UNIXTIME:$(grep cancelled $TMP_FILE | wc -l)"
```

The config Method

Apart from the output to be generated when the plugin is run without further arguments, all plugins are required to implement a config method, which is executed when the plugin is run with the config string as an argument. If we name our script muc (the abbreviation for Munich Airport) and start it from the directory where it is located, it might, for example, produce the following output:

```
$ ./muc config
graph_title Departures Munich Airport
graph_vlabel Number
graph_args --base 1000 --lower-limit 0
graph_category Departures
calling.label Calling
calling.draw AREA
boarding.label Boarding
boarding.draw STACK
departed.label Departed
departed.draw STACK
planned.label Late
planned.draw LINE2
```

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	LOGIC SUPPLY, INC. www.logicsupply.com	59
ABERDEEN, LLC www.aberdeeninc.com	35	MICROWAY, INC. www.microway.com	C4, 5
ASA COMPUTERS, INC. www.asacomputers.com	53	O'REILLY RAILSCONF www.railsconf.com	49
CARINET www.cari.net	75	O'REILLY WHERE 2.0 conferences.oreilly.com/where	49
COMMUNITY ONE communityone	15	PENGUICON www.penguicon.org	60
CORAID, INC. www.coraid.com	7	POLYWELL COMPUTERS, INC. www.polywell.com	11
COSSFEST www.cossfest.ca	61	RACKSPACE MANAGED HOSTING www.rackspace.com	C3
EMAC, INC. www.emacinc.com	67	SERVERS DIRECT www.serversdirect.com	1
EMPERORLINUX www.emperorlinux.com	27	SILICON MECHANICS www.siliconmechanics.com	21, 29
GENSTOR SYSTEMS, INC. www.genstor.com	71	SUN JAVA ONE CONF. java.sun.com/javaone/index.jsp	13
IT360 www.it360.ca	39	TECHNOLOGIC SYSTEMS www.embeddedx86.com	79
INTEROP www.interop.com	55	TECHWEB LIVE EVENTS (BLACKHAT) www.blackhat.com	33
LINUX FEST NORTHWEST www.linuxfestnorthwest.org	43	UBIQUIT NETWORKS, INC. www.ubnt.com	C2
LINUX ON WALL STREET www.linuxonwallstreet.com	34	ZT GROUP INTERNATIONAL www.ztgroup.com	3

ATTENTION ADVERTISERS

July 2009 Issue #183 Deadlines

Space Close: April 27; Material Close: May 5

Theme: Mobile Linux

BONUS DISTRIBUTIONS:

USENIX Annual Technical Conference, DC PHP, Red Hat Summit, O'Reilly Velocity

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

```
cancelled.label Cancelled
cancelled.draw LINE2
```

Each time the Munin master asks the Munin node daemon to run a plugin using the fetch command (as presented in the Telnet session above), it also executes the config method in order to find out how it should display the data in the diagram. In this example, the graph should be titled "Departures Munich Airport" (Figure 2), and the y-axis should be labeled "Number".

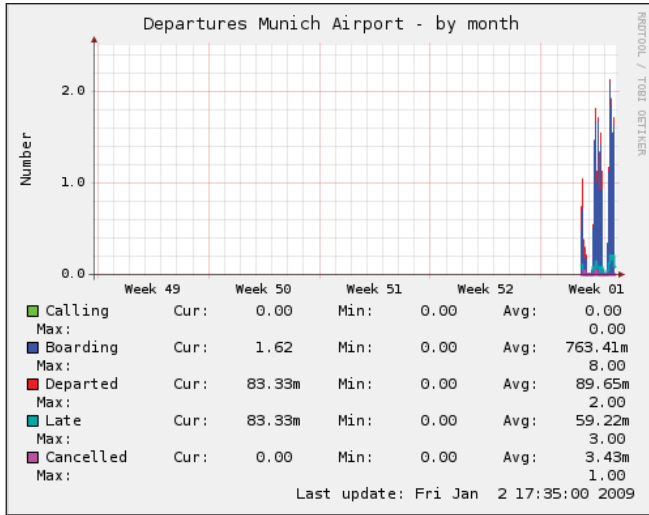


Figure 2. Departures from Munich Airport, January 2, 2009

The graph_args variable allows the plugin to forward arguments to the RRDtool graphing routine (see the rrdgraph man page). Forwarding the option --base 1000, the muc plugin ensures that a k (kilo) unit prefix as displayed in the graph equals 1000, not 1024. The -lower-limit 0 influences RRDtool's autoscaling. It makes sure that the displayed y-axis always will range at least from 0.

The graph_category tells the Munin master in which category (Figure 1) the relevant diagrams are to be displayed. This allows you to group diagrams in a logical way. The diagrams of plugins that do not specify the graph_category variable can be found in the "Other" category. The muc data will be presented in our own new category titled Departures.

How Individual Parameters Are Displayed

In addition to these variables related to the entire graph, you also can specify details for each parameter the plugin monitors, so in the case of the muc plugin, calling, boarding, departed, planned and cancelled. The <parameter>.label variable sets the legend entry for <parameter>.

<parameter>.draw specifies the type of diagram the <parameter> data is to be presented as. AREA asks the Munin master to draw the relevant curve and fill the entire area between the x-axis and data point with color. Parameter data of the drawing type STACK will be stacked on top of this basic area. This way, we sum up all flights that, in the current five-minute interval, are labeled calling, boarding and departed. The airport timetable won't correct their departure time later; they all count as dispatched at this point of time.

Flights tagged as planned and cancelled behave differently. For delayed flights, the Munich Airport authorities will issue a new departure time later. This way, the plugin will see planned flights twice: as planned within their time slot according to schedule, and as calling, boarding or departed at their actual time interval of departure. That's why we draw delayed flights in a separate line of the type LINE2. The number denotes the thickness of the line in pixels. A LINE1 line is one pixel thick; a LINE2 line is two pixels, and a LINE3 line is three pixels. Cancelled flights won't reappear in the time schedule, but as they will never depart, we also draw them as a separate line of two pixels thick.

All this plugin output is written to the standard output. The final version of our muc script is shown in Listing 1.

Listing 1. muc Munin Plugin Script

```
SITE=http://www.munich-airport.de
DEP_URL=$SITE/en/consumer/fluginfo/abflug/index.jsp?viewType=t
TMP_FILE=/tmp/.muc_flights

if test "$1" = "config"; then
    echo graph_title Departures Munich Airport
    echo graph_vlabel Number
    echo graph_args --base 1000 --lower-limit 0
    echo graph_category Departures
    echo calling.label Calling
    echo calling.draw AREA
    echo boarding.label Boarding
    echo boarding.draw STACK
    echo departed.label Departed
    echo departed.draw STACK
    echo planned.label Late
    echo planned.draw LINE2
    echo cancelled.label Cancelled
    echo cancelled.draw LINE2
else
    TIME=$(date +%H)
    MIN=$(echo "$(date +%M)/5)*5|bc)
    case $MIN in
        0) TIME=$TIME:00 ;;
        5) TIME=$TIME:05 ;;
        *) TIME=$TIME:$MIN ;;
    esac

    TMP_FILE=$TMP_FILE:$TIME
    links -dump $DEP_URL | grep $TIME > $TMP_FILE

    UNIXTIME=$(date -d$TIME:00 +%s)
    echo "calling.value $UNIXTIME:$(grep calling $TMP_FILE | wc -l)"
    echo "boarding.value $UNIXTIME:$(grep boarding $TMP_FILE | wc -l)"
    echo "departed.value $UNIXTIME:$(grep departed $TMP_FILE | wc -l)"
    echo "planned.value $UNIXTIME:$(grep planned $TMP_FILE | wc -l)"
    echo \
        "cancelled.value $UNIXTIME:$(grep cancelled $TMP_FILE | wc -l)"

    rm $TMP_FILE
fi
```

To activate the muc plugin, we simply create a symlink in the plugins directory and restart the Munin daemon:

```
# ln -s <path/to/>muc /etc/munin/plugins
# /etc/init.d/munin-node restart
```

By Telnetting from the Munin master machine to port 4949 of the node machine on which we have activated the muc plugin, we can check whether all is well. Let's see if our config method works:

```
$ telnet localhost 4949
[...]
config muc
graph_title Departures Munich Airport
graph_label Number
[...]
cancelled.label Cancelled
cancelled.draw LINE2
```

If we can do this by hand, the Munin master should generate some nice little graphs and present them via the Web, as shown in Figure 3.

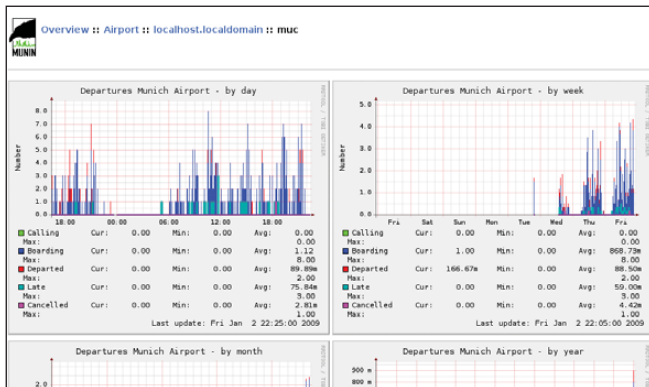


Figure 3. Munin creates daily, weekly, monthly and annual graphs. If you want to preserve them for later comparison, you can write a cron job that archives the relevant PNGs and/or RRD databases at certain points of time.

Drawbacks

The simple way to write and integrate custom plugins is one of the huge advantages of Munin—even the more complex wild-card plugins are no big deal.

Unfortunately, simple sometimes also means simplistic. Although it is possible to include Munin-generated diagrams in customized Web pages, Munin does not provide any functionality to customize the Web pages generated by the Munin master. Especially on sites with many hosts and plugins to monitor, the simple approach that combines all daily and weekly graphs in one page results in an extremely slow-loading overview page.

Another example of a rigid approach in the Munin architecture is the fixed resolution of data. Not all data changes fast enough that the five-minute interval is appropriate.

Unfortunately, a configuration option for individual plugins remains an item on the wish list. On the other hand, Munin

allows quite powerful configuration for plugins by means of environment variables set in the `/etc/munin/plugin-conf.d/` directory.

Because the documentation could be more extensive, and because the code isn't well commented, the English and German users' mailing lists remain helpful resources. ■

Patricia Jung (trish+lj@trish.de) works as a freelance technical writer, editor, community manager and system administrator specializing in open-source topics from Munich/Germany.

Resources

Munin: munin.projects.linpro.no

Integration with Nagios:

munin.projects.linpro.no/wiki/HowToContactNagios

RRDtool: oss.oetiker.ch/rrdtool


Munin Exchange: muninexchange.projects.linpro.no

rrdgraph man page:

oss.oetiker.ch/rrdtool/doc/rrdgraph.en.html

System on Module Internet Appliance Engine

- Atmel ARM9 400MHz CPU
- 6 Serial Ports & 2 SPIs
- Up to 60 Digital GPIOs
- 10/100 BaseT Ethernet
- SODIMM Bus Expansion
- SSC/I2S Audio Interface
- Linux with Eclipse IDE
- Real Time Clock Calendar
- SD/MMC Flash Card Interface
- Up to 512 MB Flash & 64 MB RAM
- 2 USB 2.0 Host Ports & 1 Device Port
- 4 10-Bit A/Ds & 6 16-Bit Timer/Counters



The SoM-9G20 uses the same small SODIMM form-factor utilized by other EMAC SoM modules, and is the ideal processor engine for your next design. All of the ARM9 processor core is included on this tiny board including: Flash, Memory, Serial Ports, Ethernet, I2S Audio Interface, PWMs, Timer/Counters, A/D, digital I/O lines, Clock/Calendar, and more. Like other modules in EMAC's SoM product line, the SoM-9G20 is designed to plug into a custom or off-the-shelf Carrier board containing all the connectors and any additional I/O components that may be required. The SoM approach provides the flexibility of a fully customized product at a greatly reduced cost. Single unit pricing starts at \$150.

<http://www.emacinc.com/som/som9G20.htm>

Since 1985
OVER
24
YEARS OF
SINGLE BOARD
SOLUTIONS

EMAC, inc.

EQUIPMENT MONITOR AND CONTROL

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

Solaris-Zones: Linux IT Marbles Get a New Bag

Solaris-Zones provides the ability to run Linux and Solaris on the same machine without all the overhead of full virtualization. VICTOR BURNS

Seldom is a data center asked to do less. More often, it's asked to do more with less—fewer computers and less power consumption. One significant industry discussion for the past few years has been regarding a reduction in the number of physical servers and an increase in the application-to-server ratio to maximize server utilization. Often, this increase is done via virtualization.

At Texas Instruments (TI), we have numerous data centers and design environments that thrive on the use of Linux and Solaris. Typically, each OS is installed on individual systems stacked high and aligned in rows throughout the data center. Linux applications run on Linux; Solaris applications run on Solaris.

Recently, a new virtualization solution has emerged that enables IT professionals to combine Linux and Solaris together within one physical environment. This solution reduces the number of physical systems in the computer environment and increases server work with greater efficiency.

One of the Solaris virtualization environments is called Solaris-Zones (also known as Solaris-Containers). Through the development of Open-Solaris, Solaris-Zones has been expanded to support zone branding. Solaris-Zones now enables the creation of "lx" branded zones. The lx branded zone supports the installation and execution of the Linux OS and its applications. When lx branded zones are used in conjunction with the ZFS (Zeta-byte File System), Linux environments are able to do more, faster.

Linux always has been about technical developers and enthusiasts doing whatever moves them. The security of Solaris-Zones combined with the power of Linux opens a huge new frontier of development freedom—from the enterprise environment to the single desktop. With Solaris-Zones, it's easy to define, create, install and execute Linux (lx) branded zones.

This article introduces lx branded zones and presents the necessary tools for each step of the zone management process. Readers should have some understanding of a chroot environment and the basic concepts of virtual machines (VMs) and the features they provide. Knowledge of these concepts is not required, but it will help in conveying what a zone is and create a better platform for understanding.

Zones Simplified

So, what is a zone? A zone provides security and virtualization in a unique way. The Solaris-Zone has its own filesystem with a root directory, system files and so on, like that of the primary environment of the physical system. The private root filesystem, one per zone, gives it the ability to be fully configurable

and flexible. A zone provides nearly the same experience as the main OS. In this regard, it is like a VM without the VM hardware emulation layer.

The zone is provided with an operating environment but without a private dedicated kernel. The lack of a dedicated kernel is a huge performance enhancement—when you experience the boot process, you will see how fast it is compared to a normal boot. User and administrator experience within the zone is very similar to that of a full VM in flexibility, but like a chroot environment, it sheds the overhead of a full VM.

It is important to understand that a zone is not a full virtual machine in the sense that you would see with Xen or VMware or VirtualBox. A zone is an emulation layer, more akin to Wine perhaps, but at a more fundamental level. This, for example, means that an lx branded zone does not contain its own Linux kernel; rather, the kernel calls are redirected by the zone's emulation layer to the underlying Solaris kernel.

The zone provides security through isolation. Each zone has its own root account and password. The superuser within a zone has no special privileges to gain access to objects outside the zone. No account has rights to exit the zone or examine processes and files outside the zone. Advanced resource management is enabled when control of memory and CPU resources by zone is important. Resource management keeps zones from being harmed by others, including but not limited to CPU and memory starvation.

Note: the primary Solaris OS and the physical platform on which it executes are also known as a zone. It is defined as the global zone and continues to look and feel as it always has. All other zones are created from the global zone. Created zones

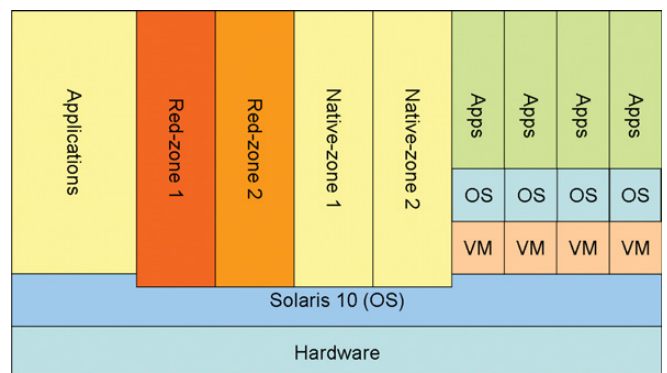


Figure 1. Relationships between Zones and VMs

are called sub or non-global zones. Non-global zones cannot create zones within themselves. Figure 1 illustrates the relationship between the global zone, non-global zones and possible VMs.

Zone Branding

Solaris-Zones became available with the release of Solaris 10 (later Open-Solaris). With these early releases, only a “native” Solaris zone could be defined, installed and executed. With the August 2007 release, Solaris-Zones includes support for zone branding to allow Linux installation and execution. By default, a zone is defined as native, unless it’s defined explicitly as a Linux (lx) branded zone. Once a zone is branded lx, only Linux can be installed into that zone.

Commands

The zone experience is defined by a simple command set. Each command is used to manage one of the logical divisions of the zone maintenance process. The primary divisions of zone administration are define, install and execute. The zone experience is very simple; it involves only a few commands. Two of the commands provide support for the definition, installation and setup of zones, and the other two are used for a running zone:

- zonecfg: define a zone (metadata only).
- zoneadm: install/uninstall, boot and query.
- zlogin: log in to a zone or connect to its console.
- zonename: prints the name of the zone executed within.

Define a Zone

Use the zonecfg command to define a zone. Although it is possible to define a zone without networking, all examples presented here define zones with networking. Listing 1 shows how to define a network interface for use by an lx branded zone. With zonecfg, you can create a minimal zone definition, set the zone’s name, set its installation path and type and include a network interface. A minimum definition requires only the branding, zone name and the installation path. The zonecfg command must be executed as the superuser. In the examples here, the shell prompt is used to illustrate from which zone a command is run. The initial example below indicates the shell is within the global zone and ready to “define” a non-global zone by the use of the zonecfg command.

Note: ZFS (denoted or hinted at by path names) is used for performance; however, it is not required. Feel free to use any appropriate directory path to build one or more zones.

Adjust the paths accordingly to match your local environment. Items to consider are zonepath and network values. Change these to match available storage, local network requirements and available network interface. The first command shows that execution is in the global zone. The zonecfg command defines the name of the zone, the installation path and network attributes. The final command lists all configured and running zones. Once a zone is defined, use the zonecfg

Listing 1. Defining an lx Zone

```
# List the name of the current zone
g-zone# zonename
global

# Start the zone definition action and define it as "lx"
# SUNWlx is the Sun provided "lx" zone template.
g-zone# zonecfg -z red-zone
red-zone: No such zone configured
Use 'create' to begin a new zone configuration.
zonecfg:red-zone> create -t SUNWlx
zonecfg:red-zone> set zonepath=/zpool01/zones/red-zone
zonecfg:red-zone> add net
zonecfg:red-zone:net> set address=192.168.1.10
zonecfg:red-zone:net> set physical=e1000g0
zonecfg:red-zone:net> end
zonecfg:red-zone> commit # (redundant)
zonecfg:red-zone> exit

# List defined(configured) and running zones
g-zone# zoneadm list -cv
```

ID	NAME	STATUS	PATH	BRAND	IP
0	global	running	/	native	shared
-	red-zone	configured	/zpool01/zones/red-zone	lx	shared

Listing 2. Changing Properties of a Defined Zone

```
# We determined the network address is incorrect and that it
# needs to be updated and we want a different storage location
g-zone# zonecfg -z red-zone
zonecfg:red-zone> set zonepath=/zpool01/zone/red-zone-x
zonecfg:red-zone> select net address=192.168.1.10
zonecfg:red-zone:net> set address=192.168.2.10
zonecfg:red-zone:net> end
zonecfg:red-zone> exit

# Use zonecfg's "info" sub-command to list the zone's definition
g-zone# zonecfg -z red-zone info
zonename: red-zone
zonepath: /zpool01/zone/red-zone-x
brand: lx
autoboot: false
bootargs:
pool:
limitpriv:
scheduling-class:
ip-type: shared
net:
    address: 192.168.2.10
    physical: e1000g0

# The zone definition can be deleted
g-zone# zonecfg -z red-zone delete -F
```

command to update or delete a zone configuration.

Note that not all properties can be updated or added after a zone has been installed. Generally, properties with this restriction are ones related to native zone definitions, not lx branded zones. For properties that can be changed after a zone is installed, the zone should be in a halted state or rebooted to make the change active.

The first example shows the red-zone as configured. This

Listing 3. Install a Zone

```
# List the zone definition (from Listing 1)
g-zone# zonecfg -z red-zone info
zonename: red-zone
zonepath: /zpool01/zones/red-zone
brand: lx
autoboot: false
bootargs:
pool:
limitpriv:
scheduling-class:
ip-type: shared
net:
    address: 192.168.1.10
    physical: e1000g0

# The ISO files:
g-zone# ls /zpool01/rh-media/rhel30-AS-U6/
rhel-3-u6-i386-as-disc1.iso  rhel-3-u6-i386-as-disc3.iso
rhel-3-u6-i386-as-disc2.iso  rhel-3-u6-i386-as-disc4.iso

# Install linux "desktop" distribution
g-zone# zoneadm -z red-zone install -d
  =>/zpool01/rh-media/rhel30-AS-U6 desktop
Installing distribution 'Red Hat Enterprise Linux 3 Update 6'...
Installing cluster 'desktop'
Installing miniroot for zone 'red-zone'.
Attempting to locate 30 packages...

Installing 1 miniroot package...

Attempting to locate 29 packages...

Installing 29 miniroot packages...
...
Installation of zone 'red-zone' completed successfully.

Details saved to log file:
  "/zpool01/zones/red-zone/root/var/log/red-zone.install.1119.log"

# Query the zones status (changed to installed)
g-zone# zoneadm list -cv
ID NAME      STATUS   PATH                               BRAND  IP
0  global    running /                                     native shared
-  red-zone  installed /zpool01/zones/red-zone          lx     shared
```

means it is defined only (metadata created and saved). Two properties in the example can be used to illustrate updating properties of an already-defined zone: zonepath and the network attributes. Each of them can be changed while the zone is halted (not running). If a zone has been installed and the zonepath is changed, the operator is required to move the physical location of the old zonepath to the location of the new zonepath manually. In the next example (Listing 2), the directory red-zone needs to be renamed to red-zone-x under the /zpool01/zones directory to complete the property update.

Install a Defined Zone

We now have a defined zone. Use the zoneadm command to complete the OS installation into the zone named red-zone. The sub functions of zoneadm are related to the execution status of a zone. The install process of an lx branded zone requires Linux media. The media can be provided in a physical form and loaded into the system's CD-ROM drive, or you can use the "green" method and provide the image as one or more ISO files.

Boot the Zone

Once the zone installation is complete, it's time to boot it. Create two shells, and run the commands shown in Listing 4. Connect to the zone console first, then boot the zone in the second shell to get the full console experience (it's very fast, you'll not want to miss it). The example zlogin connects to the

Listing 4. Boot a Zone

```
# Connect to the red-zone console and watch it boot
g-zone(1)# zlogin C -e '#' red-zone

# Boot the red-zone in a second window
g-zone(2)# zoneadm -z red-zone boot

# **** example console output from first window ****
[Connected to zone 'red-zone' console]

[NOTICE: Zone booting up]
INIT: version 2.85 booting
        Welcome to Red Hat Enterprise Linux AS
        Press 'I' to enter interactive startup.
Configuring kernel parameters: [ OK ]
Setting clock : Mon Sep 29 20:20:56 EDT 2008 [ OK ]
Setting hostname red-zone: [ OK ]
...
Starting atd: [ OK ]
Starting firstboot: [ OK ]
Rotating KDC list [ OK ]

Red Hat Enterprise Linux AS release 3 (Taroon Update 6)
Kernel 2.4.21 on an i686

red-zone login:
```


zone's console device and configures the escape (exit the zlogin) as the "#." (pound sign then period) key sequence. This key sequence should be unique and avoid issues that the default sequence of "~." (tilde then period) can cause when connectivity to the global zone is remote.

Login

A non-global zone has nearly the same abilities as the global zone to provide services: login connections are not limited to text or console logins. The use of zlogin with no options (only the zone name) connects to the zone without a console, which creates a tty and invokes login. Any active zone service also can be used, such as XDM, SSH and FTP, to allow other forms of login.

Listing 5. Log In to a Zone

```
# Connect to zone's console
g-zone# zlogin -C -e '#' red-zone

# Login without a console (must always be root)
g-zone# zlogin red-zone
...
r-zone# zonename
red-zone
```

Zone Administration

We now have a zone defined, installed and running. The examples presented here illustrate some of the administrative tasks associated with zones: reboot, shutdown, halt and deletion of an lx branded zone. Pay close attention to the shell prompts to identify the zone in which each command is run.

Automated Zone Creation (zonetool.pl)

The zone creation steps are straightforward and simple. The process may take only a few steps, but they are manual and error-prone. The zonetool.pl utility (see Resources) automates the zone creation process and includes detailed POD documentation. Run zonetool.pl without arguments or with the --help option to display usage details. Listing 7 shows an example of using zonetool.pl.

Closing Thoughts and Ideas

With relatively small amounts of disk and memory resources, a single physical server can host hundreds of zones. Each zone is usable by any number of users, and a single-user zone provides extreme flexibility. A single user can create more than one zone to test both server and client environments, and the applications will believe they are on unique physical hosts. A zone user may have use of the zone's unique root password or unfettered sudo access within that zone without concern for security and stability of the global zone and other non-global zones.

The lx branded zone does have its limitations. Much of the zone's power comes from securely shared resources with the global zone. The zone shares a kernel with the global zone



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.

Listing 6. Zone Administration

```
# By default a zone will *not* auto boot when the system boots
g-zone# zonecfg -z red-zone set autoboot=true

# Reboot/stop a zone gracefully
g-zone# zlogin -C -e '#' red-zone

# Login, then use a normal Linux command (examples)
r-zone# reboot
r-zone# shutdown -r now
r-zone# telinit 6
r-zone# telinit 0 (once halted boot it from global zone)

# Bring down a zone hard (stop all execution now!)
g-zone# zoneadm -z red-zone halt

# Same as halt and boot (hard stop then boot)
g-zone# zoneadm -z red-zone reboot

# Uninstall the OS from a zone (still defined)
g-zone# zoneadm -z red-zone halt
g-zone# zoneadm -z red-zone uninstall [ -F ]

# Clone an installed zone's OS into another defined zone
# (when the source is on ZFS the clone takes seconds to
# install the new zone, otherwise may take many minutes!)
g-zone# zonecfg -z red-eye create -t red-zone
g-zone# zonecfg -z red-eye set zonepath=/zpool01/zones/red-eye
g-zone# zonecfg -z red-eye
zonecfg:red-eye> select net address=192.168.1.10
zonecfg:red-eye:net> set address=192.168.1.11
zonecfg:red-eye:net> end
zonecfg:red-eye> exit

g-zone# zoneadm list -cv
ID NAME STATUS PATH BRAND IP
0 global running / native shared
- red-zone installed /zpool01/zones/red-zone lx shared
- red-eye configured /zpool01/zones/red-eye lx shared

g-zone# zoneadm -z red-eye clone red-zone
Cloning zonepath /zpool01/zones/red-zone...

g-zone# zoneadm list -cv
ID NAME STATUS PATH BRAND IP
0 global running / native shared
- red-zone installed /zpool01/zones/red-zone lx shared
- red-eye installed /zpool01/zones/red-eye lx shared

# Remove a zone (No undo)
g-zone# zoneadm -z red-eye halt
g-zone# zoneadm -z red-eye uninstall -F
g-zone# zonecfg -z red-eye delete -F
```

Listing 7. Zone Definition Using zonetool.pl

```
$ sudo zonetool.pl \
    [ --debug ] \
    --create \
    --zonename green-zone \
    --brand SUNWlx[,server ] \
    --zonepath /zpool01/zones \
    [ --autoboot ( true | false ) ] \
    [ --clone red-zone | --media /path/to/ISOs ] \
    --network 'e1000g0=192.168.1.10'
```

and, therefore, places limits on kernel modules and drivers. Because zones are not full VMs, the Linux distributions that can be installed in an lx branded zone are limited. Support for other Linux releases is possible, and further interest in this technology will inspire continued development and support for additional Linux distributions. Review the Resources section of this article for more information on this and related topics. ■

Victor Burns is married to his awesome wife Lisa of 25 years, and they have six children. Victor has been employed at Texas Instruments for 24 years. He enjoys working with the Boy Scouts and is an Eagle Scout from his youth. He thanks his family for its support.

Resources

Location of brandz installation extensions and other helpful items: opensolaris.org/os/community/brandz/downloads and opensolaris.org/os/community/brandz/todo/linux_2_6 (kernel 2.6 support).

Solaris lx branded zone information: docs.sun.com/app/docs/doc/819-2450/gchhk?a=view (everything), docs.sun.com/app/docs/doc/819-2450/gcwwm?a=view (planning), docs.sun.com/app/docs/doc/819-2450/gdajn?a=view (config), docs.sun.com/app/docs/doc/819-2450/gdbki?a=view (about...), docs.sun.com/app/docs/doc/819-2450/gdduh?a=view (install, boot, and so forth) and docs.sun.com/app/docs/doc/819-2450/gdqnv?a=view (moving and migrating).

zonetool.pl (Listing 8): ftp.linuxjournal.com/pub/lj/listings/issue180/10268.tgz.

zonetool.pl was inspired by an early version by M. Kiefer: www.sun.com/bigadmin/scripts/submittedScripts/zonetool.txt.

Dtrace information: www.sun.com/bigadmin/content/dtrace, wikis.sun.com/display/DTrace/Documentation and docs.sun.com/app/docs/doc/819-5488?l=en.

ZFS information: docs.sun.com/app/docs/doc/817-2271?l=en.

3-D Graphics Programming with Irrlicht

Stone obelisks, a purple spinning cube and a Babe in Black—all this in less than 100 lines of code! MIKE DIEHL

There's something about 3-D graphics that just draws you in. Even though I have a degree in Mathematics, I've always had the impression that programming in 3-D would be difficult. I've recently discovered though, that it's really not very hard. In fact, it's almost easy and a lot of fun—thanks to the Irrlicht 3-D graphics engine.

The Irrlicht 3-D graphics engine is written in C++ and allows you to get impressive results without a whole lot of code, as you'll see later in this article. With Irrlicht, you can write programs that will run under Linux or Windows and take advantage of OpenGL or DirectX. Irrlicht directly supports 3-D models in various formats, including Maya (.obj), COLLADA (.dae), Quake 3 levels (.bsp), Quake 2 models (.md2) and Microsoft DirectX (.X), among others. This means there are plenty of ready-made models available for download on the Internet that can be used with Irrlicht. Also, many tools are available for creating models and textures to use with Irrlicht.

When I was evaluating some other 3-D engines, I chose Irrlicht because it seemed to be the easiest to wrap my head around, while at the same time, it had all the features I wanted. Irrlicht supports both mesh-based animation as well as a skeletal animation system. With Irrlicht, materials can be layered to produce stunning effects. And most important, Irrlicht is extremely well documented with tutorials on-line as well as a very responsive on-line forum. Oh, and it's free. And, it's open source.

Listing 1 shows a sample program I wrote in order to demonstrate a few of Irrlicht's features.

The first 18 lines of code are pretty easy to follow. They include the `irrlicht.h` header file, which contains all the declarations I'll need. Then, I define a few namespaces and variables for later use in the program. The main function begins on line 19.

On line 23, I ask Irrlicht to set up my display window. Here, I tell it to use the OpenGL render engine and to use a 640x480 display resolution. I've included comments in lines 20–22 that show what values to use to select from the various other render engines that Irrlicht supports. In line 26, I check to make sure the call to `createDevice()` was successful. If it wasn't successful, it's Game Over, literally.

Lines 27 and 28 initialize a few objects that I'll use throughout the rest of the code. The driver object allows me to change various aspects of how the window is rendered; I'll use this object in the next code block. The `smgr` object is the scene manager object and is the object I use to add objects to my scene, such as cameras, lights and other objects.

In lines 29–35, I set up what's known as a skybox. A skybox is exactly what it sounds like. Imagine a giant box that is set

down over a scene, each face of the box having a different mural on it. So, if you were to look to the west, you would see the mural on the western face of that skybox. And, if that mural were a picture of a sunset, it would present the illusion that you were looking at a real sunset. In my example here, I use the skybox textures that came with the Irrlicht tutorials.

A common mistake that first-time Irrlicht users make is building up their scene by adding all kinds of models and various types of objects, but when they go to display their creation, they don't see anything but black. You can't see anything without light. I add a light object as well as some ambient light in lines 37–39.

At this point in the code, I get my first introduction to what's known as a vector. A vector is simply an object that has more than one numerical component. In this case, it's a `vector3df` object, which simply means it's composed of three floating-point components. You can think of these components as X, Y and Z, or perhaps up/down, left/right and forward/backward. Essentially, a vector allows you to store a location in 3-D space. The `SColor` vector in line 39 also has three elements. In this

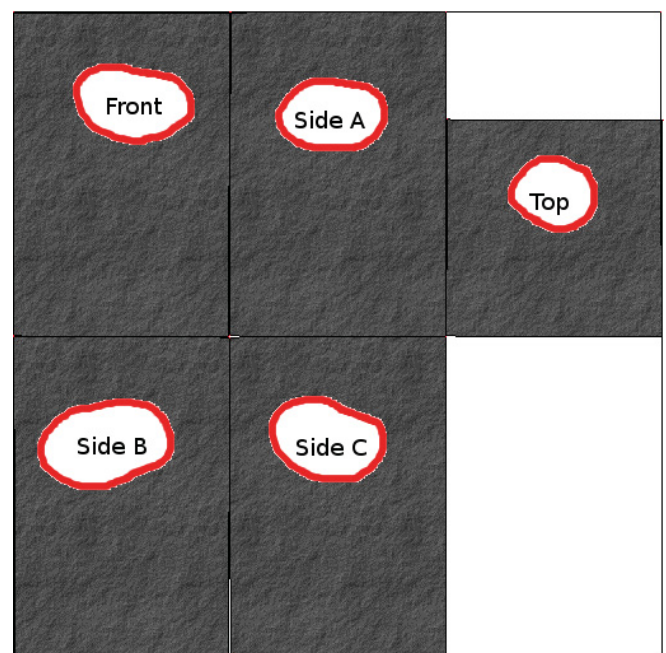


Figure 1. 3-D Texture Represented in Two Dimensions

Listing 1. Irrlicht Sample Program

```

1  #include <irrlicht/irrlicht.h>
2  #include "unistd.h"

3  using namespace irr;
4  using namespace irr::core;
5  using namespace irr::video;
6  using namespace std;

7  IrrlichtDevice*      device;
8  video::IVideoDriver* driver;
9  scene::ISceneManager* smgr;
10 scene::ICameraSceneNode* camera;

11 scene::IAnimatedMesh* ground;
12 scene::IMeshSceneNode* ground_node;

13 scene::IAnimatedMesh* house;
14 scene::IMeshSceneNode* house_node;

15 scene::IAnimatedMesh* avatar;
16 scene::IAnimatedMeshSceneNode* avatar_node;

17 video::SMaterial      material;
18 scene::ISceneNode*   cube;

19 int  main () {

20     //video::EDT_SOFTWARE
21     //video::EDT_NULL
22     //video::EDT_OPENGL,

23     device=createDevice(video::EDT_OPENGL,
24     dimension2d<32>(640,480),16,false,true);
25
26     if (device == 0) return(1);

27     driver = device->getVideoDriver();
28     smgr = device->getSceneManager();

29     smgr->addSkyBoxSceneNode(
30         driver->getTexture("./graph/irrlicht2_up.jpg"),
31         driver->getTexture("./graph/irrlicht2_dn.jpg"),
32         driver->getTexture("./graph/irrlicht2_lf.jpg"),
33         driver->getTexture("./graph/irrlicht2_rt.jpg"),
34         driver->getTexture("./graph/irrlicht2_ft.jpg"),
35         driver->getTexture("./graph/irrlicht2_bk.jpg"));
36
37     smgr->addLightSceneNode(0, vector3df(0, 100, 0),
38         video::SColorf(1.0f, 1.0f, 1.0f), 1000.0f, -1);
39     smgr->setAmbientLight(video::SColorf(255.0,255.0,255.0));

40     camera = smgr->addCameraSceneNodeFPS(0,30.0f,90.0f,-1,
41         0,0,false,0.0f);
42     camera->setPosition(vector3df(30,10,30));

43     ground = smgr->getMesh("./graph/grass.obj");
44     ground_node = smgr->addMeshSceneNode(ground);
45     ground_node->setScale(vector3df(1000,1,1000));
46     ground_node->setMaterialFlag(EMF_LIGHTING, false);

47     material.setTexture(0,
48         driver->getTexture("./graph/building.tga"));
49     house = smgr->getMesh("./graph/building.obj");

50     for (int i=0; i<5; i++) {
51         house_node = smgr->addMeshSceneNode(house);
52         house_node->setScale(vector3df(.5,.5,.5));
53         house_node->setPosition(vector3df(30*i+5,0,-30));
54         house_node->getMaterial(0) = material;
55         house_node->setRotation(vector3df(0,90,0));
56     }

57     material.setTexture(0,
58         driver->getTexture("./graph/sydney.bmp"));
59     avatar = smgr->getMesh("./graph/sydney.md2");
60     avatar_node = smgr->addAnimatedMeshSceneNode(avatar);
61     avatar_node->setScale(vector3df(.1,.1,.1));
62     avatar_node->setPosition(vector3df(5,2.5,5));
63     avatar_node->setRotation(vector3df(0,270,0));
64     avatar_node->getMaterial(0) = material;

65     cube = smgr->addCubeSceneNode(1.0f, 0, -1,
66         vector3df(10, 2, 10),
67         vector3df(45.0, 0, 0),
68         vector3df(1.0f, 1.0f, 1.0f));
69     cube->setMaterialTexture(0,
70         driver->getTexture("graph/purple.jpg"));

71     cube->addAnimator(
72         smgr->createRotationAnimator(vector3df(1,.5,.25)));

73     while (device->run()) {
74         driver->beginScene(true,true,
75             video::SColor(255,100,101,140));
76         smgr->drawAll();
77         driver->endScene();
78     }

79     driver->drop();

80     return(0);
81 }

```

case, it's safe to think Red, Green and Blue.

Lines 40 and 41 are probably the most important. Without them, I still would not see anything, nor would I be able to "walk around" in my scene. In line 40, I add a First Person Shooter (FPS) camera to the scene. It's this camera that determines what I see. It's also this camera that I move around with the arrow keys and mouse. The FPS camera is my eyes into the game. The Irrlicht engine supports various other camera types, but the FPS camera is the most intuitive, because it mimics the FPS games everyone's familiar with. In line 41, I position the camera at a location described by the vector, (30,10,30).

Lines 42–45 are where I add my first mesh to the scene. Think of a mesh as just a bunch of triangles and rectangles that are put together to form the shape of an object. In this case, I'm adding a simple rectangle shape to form the ground in my demonstration. First, I call `getMesh()` to read a mesh from an external file. Then, I call `addMeshSceneNode()` to convert that mesh into a local representation and add it to my scene. This function returns an object that gives me access to that representation. Using this object allows me to use the `setPosition()` and `setScale()` methods to move the mesh around and set its size in my scene. Finally, I use the `setMaterial()` method to tell Irrlicht that this object does not emit light on its own.

At this point, I have a sky, some light to see by, a camera to see with and some ground to stand on. But, it gets better.

I put in a few background objects in lines 46–54. In this block of code, I create my first material by reading in an external texture file. This material then will be applied to the meshes as I add them to my scene. On line 47, I read in the mesh that eventually will be used to add a row of stone "houses" to my scene. Inside the loop, I add them to the scene, scale them, position them in a row, and turn them around a bit.

Finally, on line 52, I apply the material that I created on line 46. Figure 1 shows what's inside `building.tga`—it's what might happen if you thought of the house as simply a box and "unfolded" it so all of its sides fit flat on a piece of paper. I then added a slate texture and a label to each face. When I apply this material to the mesh in line 52, the faces from `building.tga` are wrapped around the model to form an object that appears to be made of stone. This process is known as UV mapping.

Lines 55–61 expose as much complexity as you're going to see in this short example. Here, I reuse the material variable, which is probably bad form, but this is only meant to be a quick demonstration. This time, I'm reading in a UV

mapping that is considerably more complex than the box I created for the houses earlier. This material is used as a skin for the `sydney.md2` Quake model. In line 57, you can see that this is an animated mesh, which is different from the meshes discussed so far. An animated mesh contains several meshes that can be used, in turn, to create animations. In this case, Sydney has various death-scene animations. She also has a running animation. Sometimes, at one point, I swear she's doing the Macarena! The rest of the code block is devoted to scaling, positioning and rotating the model to our liking.

Now things get a little psychedelic. In lines 62–66, I create a cube that appears to be floating just above the ground. I also apply a purple skin to it. Sure, purple floating cubes are one thing, but on line 67, I make it rotate in space. To add to the visual effect, I specify that the cube revolves once per second around the X axis while revolving around the Y axis twice a second, and finally, it revolves around the Z axis three times a second. The result is a cube that floats in space and spins around in an apparently random fashion.

Lines 68–72 are the main run loop. The `run()` method returns true until users press the Esc key, indicating they want to end the game. If this game needed to move objects around,



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




Figure 2. The Completed Scene

such as flying missiles or attacking bad guys, those game updates would take place between the call to `beginScene()` and the call to `drawAll()`.

Finally, when users press the Esc key, I release some resources in line 73, and the program returns to the operating system.

I can compile the program with a command that resembles this one:

```
g++ ./lj.cpp -lIrrlicht -lGL -lXxf86vm -lXext -lX11
➔-lenet -ljpeg -lpng -o game
```

See Figure 2.

So there you have it. The example here shows building a simple scene, adding a moving character and a spinning cube. You even can walk around and explore this simple world, or fly around and explore it from above, or below—all this in less than 100 lines of code!

As powerful as Irrlicht is, it's not without its weaknesses. I've not had much success with Irrlicht's support of materials other than UV-mapped materials. I desperately tried to get the ground to resemble actual grass, but I couldn't seem to get it to work. On the other hand, UV mapping a complex model is a daunting challenge. I also noticed some of the tools used to create or export the models behaved strangely. Sometimes the resulting models would look okay. Other times, they'd need to be rotated or scaled in order to look right. Of course, most of

these problems are problems with the 3-D modeling tools used to create content for Irrlicht, and not with Irrlicht itself.

I've also discovered that writing a 3-D game is more about artwork than code-work. This simple demonstration has all the major elements of an FPS game. But the scene is still quite simple and not very realistic. However, by merely changing the models and textures, this scene could be made to look like a row of realistic houses with doors and windows, perhaps a street and sidewalk complete with grassy lawns and a newspaper sitting on the porch—no code changes required. I simply could have used an existing Quake or Doom level instead, but I'm kind of tired of the Gothic atmosphere of most of those games. I'd like to see a new crop of brighter, more familiar-looking FPS or MPORPG games.

While researching this article, I examined some of the competing 3-D graphics libraries. Ogre seems to be the leading contender in my opinion. From reading the user's manual, I formed the impres-

sion that Ogre had a more intuitive API, but that I'd have to write much more code to get the same results that I do with Irrlicht. I also was put off by the fact that Ogre supports only a single mesh format, although exporters are available for converting other formats.

As you might have guessed, I'm writing a 3-D game using Irrlicht. However, I started this project as an excuse to learn C++. When I started, I really thought the stumbling block would be writing the code needed to make the game functional. I've discovered that the hard part of writing any 3-D game is in the artwork. Creating compelling scenes and realistic landscapes with trees and shrubs is hard. Coding is relatively easy, thanks to advanced libraries like Irrlicht. The example in this article doesn't even begin to scratch the surface of what Irrlicht can do. Indeed, I've not even begun to scratch the surface in my programming efforts. ■

Mike Diehl is a self-employed computer consultant and lives in Albuquerque, New Mexico, with his wife and three sons. He can be reached at mdiehl@diehlnet.com.

Resources

Irrlicht Home Page: irrlicht.sourceforge.net

Description of Quake 2 model files (md2):
tfc.duke.free.fr/coding/md2-specs-en.html

Irrlicht Support Forum:
irrlicht.sourceforge.net/phpBB2/index.php

Did you know Linux Journal maintains a mailing list where list members discuss all things Linux? Join LJ's linux-list today:
<http://lists2.linuxjournal.com/mailman/listinfo/linux-list>.

Mutt vs. Thunderbird

Kyle and Bill face off on mail client software.

In our argument...I mean, in our last column, we discussed the pros and cons of owning a large laptop versus a small one. This month, we delve into something a little more personal: choice of mail client software and why one would want to use a text-based mail client over one that's graphical.

BILL: I'm a Thunderbird user, and Kyle prefers the retro-exoticism of Mutt. I'll let Kyle kick off the conversation, as I'm sure he'll have something pithy to say about what an amazing and cool mail program Mutt is.

KYLE: For starters, let's talk about startup times. Because Mutt is so lightweight, I don't even notice the startup times. For instance, I just timed how long it took me to start Mutt and have it open the INBOX and then exit as fast as I could:

```
greenfly@minimus:~$ time mutt
Mailbox is unchanged.
real    0m0.782s
user    0m0.096s
sys     0m0.036s
```

That's right, 0.8 seconds. I could have been even faster, but my reflexes are slowing in my old age. In fact, Mutt starts up so fast, that unlike other programs on my desktop—*cough*, Firefox *cough*—I don't have to leave it open. I just press a hotkey and launch it when I have new e-mail.

BILL: I just timed Thunderbird, and it took about ten seconds to start up and be done checking mail. Thunderbird may take ten seconds to launch and check mail in all my folders, but it's not like ten seconds is a large amount of time. Once it's open, I leave it open until I'm done with the computer, so it's a ten-second hit, once. There's more to life than raw speed. Good luck with LDAP lookups for corporate address books using Mutt.

KYLE: Oh yes, ten seconds sure sounds a lot like 0.8 seconds. And yes, LDAP lookups. I guess that would be a problem if Mutt weren't so extensible. Because you can make its query command anything you want, you can point to any script. The Internet was kind enough to provide an LDAP query script for me, and other people have queried Gmail and other contact databases.

BILL: Whatever. You need Postfix or some other

MTA to send e-mail. What kind of ridiculousness is that? Yeah, another server process on my *laptop* to maintain and troubleshoot. That sounds like a great idea. Thunderbird can talk SMTP (and even SMTP over TLS, shocking though it may seem!) natively. Even Eudora has enough brains to talk to an SMTP server natively, and that was back in 1994. Oh, wait. You use Mutt. Hey, Kyle, 1989 called; they want their VT-100 back.

KYLE: Who needs to talk to an MTA directly when every Linux system has its own mail server? A local mail server will handle delivering my mail much better and stabler than any client could. What ever happened to "Do one thing and do it well"?

BILL: That's arguable, but beyond the scope of this column. Besides, Mutt has other things that drive me batty...like the user interface. It drives me insane. It's hyper-pedantic. Gotta mark all messages as read? Ctrl-Alt-D+*+infinity sign or something. It's so much easier to right-click on the Inbox and say "Mark all as read" in Thunderbird.


KYLE: Yeah, easier up front and slower for the rest of time. Sure Mutt uses keybindings (familiar ones if you have ever used vim), but the learning curve pays off big-time as you blow through your mail. I can launch Mutt, read through all my new mail, and close it, before Bill's Thunderbird process finishes loading.

BILL: But, then I lose a valuable caffeine break. I click on Thunderbird, have a sip of Coke Zero, and take a quick peek at the latest LOLcats while it loads—no bother at all. What about mail delivery to Mutt? It doesn't support IMAP IDLE, does it? If you're all about speed, you'd want that feature. I seem to remember being more "on top" of e-mail than you were, back when we sat across from each other.

KYLE: The only systems I've ever noticed any IMAP issues on, were ones back-ended on Exchange. Mutt always has gotten fast notifications from my Linux-based IMAP servers.

BILL: And, Exchange is perhaps the most popular mail server in corporate environments.

KYLE: What about the RAM footprint? I wanted to get some metrics for this column from top, but



KYLE RANKIN



BILL CHILDERS

when I launched it and sorted by RAM, Mutt didn't even show up on the console. I had to stretch my terminal until it filled the entire screen before I could find the Mutt process.

BILL: My Thunderbird process shows it eating, I mean, using 128MB of RAM. I don't mind, I have 4 gigs in this thing. What about calendar integration? Thunderbird offers a plugin called Lightning to integrate my Google Calendar in one interface. Try that with Mutt.

KYLE: Oh, that's right—you have a Sport Utility Laptop with 4GB of RAM. In any case, I never understood checking your calendar with your e-mail program. It always seemed like two unrelated things—unless I guess you were born and raised on Outlook.

BILL: Here's a clue: most of the world was born and raised on Outlook.

KYLE: My magic eight ball says "Outlook not so good." Anyway, another great thing about Mutt is that I can check my e-mail from any system in the world, provided I can get access to putty (if it's Windows) or a Linux terminal to ssh into my server. Bill has to download and install Thunderbird, configure it, and then pay for another hour at the Internet café.

BILL: No, that's when I use my BlackBerry. If that's your "speed" use case, I'll put my quickdraw BlackBerry up against your mini laptop anytime, Black Bart.

KYLE: Certainly sounds like Thunderbird is great, since you have to supplement it with one or two other mail clients.

BILL: Man, my grandmother "gets" Thunderbird. I put Mutt in front of her, and I may as well have asked her to prove Ohm's Law using calculus.

KYLE: Ohm's Law, eh? Well not everyone shows such resistance (pun intended) to Mutt. Another thing Mutt got right is configurability. Just about anything you would want to configure, you can configure in Mutt. At first, I spent some time getting things just right, but ultimately, Mutt is like a finely tailored suit—you do a few tweaks here and there, and then it fits you perfectly. These days, I almost never change any settings.

BILL: Thunderbird has tons of configurability. It's a Mozilla-spawned app, not a GNOME-spawned app. And, who spends a lot of time configuring a mail client? Set it up, point it at your mailbox, and

you're done. Fire and forget. I guess pedantic people would tweak their mail settings periodically, so they'd want a pedantic client.

KYLE: I know I mentioned vi keybindings, but that bears repeating—there's nothing quite like being able to manage the entire e-mail program strictly from the keyboard and even have it do all the really heavy lifting for me. There are configurable hooks for everything. So, if I want a particular signature for my LUG e-mails and a different one for my friends, Mutt takes care of it for me.

BILL: Vi keybindings, yeah yeah. I just showed my wife, Kelly, what Mutt looks like. She said, "Mutt looks like something I'd use on the Commodore 64. Kyle uses that? What's that about?"

KYLE: That's because it probably could run on a Commodore 64. Thunderbird's RAM footprint couldn't even be *stored* on one.

BILL: I come back to your speed argument again. I just timed Thunderbird starting up, checking mail in all folders and shutting down, and it took six seconds this time. That's not "get up and go to the bathroom" time, that's barely a couple blinks.

KYLE: No, that happens if you ever want to do something like use Thunderbird with local maildirs. The entire thing goes down the crapper.

BILL: And *why* would you want to do that? Store it all on the server. IMAP for the win.

KYLE: Hey, if your laptop is always on-line, why not use Gmail? For me, I like to be able to reference archived e-mail no matter what connection I have. Mutt can handle just about any e-mail source you can throw at it.

BILL: If that's an issue, just set Thunderbird to cache e-mail locally, then. Configurability.

KYLE: Then there's the mailing-list support. It's hard to beat Mutt for managing multiple folders of mailing lists. It makes it easy to sort through threads without thinking every IMAP folder is for a mailing list.

BILL: Thunderbird has a threaded reader mode.

KYLE: Must be fun to set that every time.

BILL: It's a couple clicks. What do you have against a mouse?

KYLE: I can type at almost 100 words per minute.

I don't even know how to manage mouse speed, but I know it's slower. Back to configurability, managing PGP is a breeze with Mutt. You even can have it sign and encrypt replies to encrypted messages automatically, or have it always encrypt messages to particular people.

BILL: Managing PGP is easy with Thunderbird too—simply add the Enigmail plugin and configure it. It has all those features you mention.

KYLE: All of this support is built in. I didn't have to grab any plugins, and it still doesn't take up much RAM.

BILL: And, it still looks like a Commodore 64 application and is deliberately obtuse.

KYLE: Turn in your geek card now. All the geeks I know think looking like a Commodore 64 application is a bonus—especially on my green-on-black terminal.

BILL: Having a geek card is for geeks, man. You'll learn this when you surpass "system architect" and become an "IT Manager". And I told you, 1989 called; they want their VT-100 back.

KYLE: Ahhh, I see what happened. One step into management, and Bill started pointing and clicking at everything. Next thing you know, he'll be complaining because he can't embed his PowerPoint presentation into the body of his Thunderbird e-mail.

BILL: Hey, even *you* dug the data-center calculator spreadsheet I came up with. Fear my OpenOffice.org mojo. I seem to remember you wanting to learn the formula-fu for that one.

KYLE: Heh, I don't remember that part. I'd be too scared I'd wake up the next morning with pointy hair.

BILL: You have to have hair long enough to be pointy.

KYLE: Back on topic, man. Another great thing about Mutt is that it makes it easy to check local mail on any server

in your environment. It's much better than installing X libs, or firing up less on `/var/spool/mail/`.

BILL: Now *that* is a point I will concede, which is why I do maintain a Mutt instance on my colo'd server to access via ssh. But, I use it only about once a year. And, every time I do, I feel the pain of the interface. Reminds me of the first time I tried to use Debian's dselect program.

KYLE: Do you remember when you had too many messages in your INBOX to delete, and Thunderbird choked? I remember you asking me about my mutt-fu then too.

BILL: No, that was Exchange's IMAP support choking, not Thunderbird. I have more messages in my Zimbra mail, and it handles mass deletes just fine.

KYLE: Ahh, yet another mail program, eh? Right. All I know is how quickly Mutt churned through all of those e-mail messages.

BILL: Back to speed again. For a guy who drives a Honda, you're sure obsessed with speed.

KYLE: I'm more obsessed with efficiency than speed in both my car and my mail client. What it ultimately comes down to is efficiency in the user interface and RAM footprint, along with the fact that I can tweak just about any setting I *don't* like. E-mail is just one of the many things I deal with in a day, so I like to get in, get out and get on with my work. The fact that I can access the same mail client on any computer and even use it on each of my servers is a bonus. In conclusion, it has vi keybindings. Enough said.■

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.

Bill Childers is an IT Manager in Silicon Valley, where he lives with his wife and two children. He enjoys Linux far too much, and he probably should get more sun from time to time. In his spare time, he does work with the Gilroy Garlic Festival, but he does not smell like garlic.

7" Touch Panel Computer for embedded GUI / HMI applications



quantity 1 pricing starts at **\$449**

Powered by a
200 MHz ARM9 CPU

- Low power, Industrial Quality Design
- Mountable aluminum frame
- 64MB SDRAM (128MB opt)
- 512MB Flash w/ Debian Linux
- Programmable FPGA - 5K LUT
- 7" Color TFT-LCD Touch-Screen
- 800x480 customizable video core
- Dedicated framebuffer - 8MB RAM
- Audio codec with speaker
- Boots Linux 2.6 in about 1 second
- Unbrickable, boots from SD or NAND
- Runs X Windows GUI applications
- Runs Eclipse IDE out-of-the-box

Our engineers can
customize for your LCD

- Over 20 years in business
- Never discontinued a product
- Engineers on Tech Support
- Open Source Vision
- Custom configurations and designs w/ excellent pricing and turn-around time
- Most products ship next day

See our website for our
complete product line



We use our stuff.

visit our TS-7800 powered website at
www.embeddedARM.com
(480) 837-5200



Exchanging Exchange

On progress toward dislodging Microsoft's biggest enterprise lock-in. DOC SEARLS

More than five years ago, Linux-friendly IT executives told me (www.linuxjournal.com/article/7414) there was one Microsoft product that was not only highly entrenched but appealing as well: Exchange Server. "We can easily see our way to replacing Microsoft Office and even Microsoft Windows", one executive told me. "But we can't get along without Exchange. If you're looking for Microsoft's real lock-in with enterprise customers, Exchange is it."

That was then, but what about now?

Recently, I asked readers on the *Linux Journal* Web site how they'd replace Exchange.

Zimbra came up big:

- "Zimbra is open source (though not free software) and has a feature set that is on par with and sometimes even exceeds Exchange. If you are not running an MS-only shop, Zimbra is the best alternative."
- "Zimbra is extremely good, though not light on resources, and it requires a dedicated server to run. It does work flawlessly with free clients like Thunderbird and Evolution while having good support for Outlook. The downside is that it is not completely free (the YPL is open source, but entirely free)."
- "Zimbra, hands down. If you're looking for an open-source alternative to Exchange, one of the most important things to look at is the community. The Zimbra community on the forums is awesome, and the Zimbra staff really pay attention to what we as the community say."
- "I also have used Zimbra (free) for my other gigs, and it has worked well, but the packaging of Zimbra is not very desirable."
- "I love Zimbra, as it has a cohesive PIM suite that simply and seamlessly syncs with my mobile devices."
- "We switched to the Zimbra Open Source version, and it is working quite well. There was some initial resistance from users, but that has largely dissipated over a couple months of use."
- "Zimbra (Network Edition, not the free/OSS edition) gives full Outlook integration and prob-

ably works just as good as Exchange, functionality-wise. Unfortunately, for my customers (very small business), Zimbra's pricing is not up to par, and compared to Windows SBS pricing, even way too expensive. But, I guess for enterprise environments, it's a very good option."

- "Zimbra is probably the best open-source enterprise e-mail/collaboration product on the market."

There was some support for Google Apps:

- "Google Apps does not have the high degree of client seamlessness that Zimbra has: Web, complete Outlook (PIM, Filters, GAL, etc.), Zimbra Desktop, iPhone/WAP, OTA ActiveSync and so on."
- "We're small (20-person software company), but we dropped Exchange about two years ago in favor of Google Apps. We lost seamless contact and calendar sync on some mobile phones in the office, but we gained a bunch of time from our IT administrator (that was otherwise spent administering Exchange). We gained a great deal of stability and reliability, and we gained all the benefits of Google Mail."
- "Google uses Gmail. Doesn't everybody?"
- "...just throw away Exchange and use Google Apps/Docs/Mail, as many businesses do. Google Apps are so much better!"
- "Our university, including our medical school and hospital, went with Google, and it has been a big success."

IBM Lotus Domino/Notes was endorsed for big enterprises, but not without respect for Exchange:

- "We've installed IBM Lotus Domino/Notes at a dozen small (10-100 users) businesses. Its versatility is awesome. It's so much more than e-mail. It's rock solid. We run servers on SUSE, Red Hat and Windows. We run clients on Windows, Mac and, yes, Linux! And the pricing is so very affordable."
- "I used to work for a multinational (100,000+ employees) consumer goods company who used Notes/Domino, then was moved to Big

Blue who also used Notes/Domino. I now work for a much smaller bank (6,000 employees) that uses Outlook/Exchange. Experience has shown that larger companies with massive user bases tend to use Domino, as it just scales better than Exchange used to. Microsoft has put a lot of effort into making sure Exchange is more robust and that it scales better...Having used both, my preference is for the Notes/Domino combination, as I miss having my workspace."

Citadel got props:

- "...you really want to take a look at Citadel. Although it looks nothing like Exchange, it solves many of the same problems (e-mail, calendars, address books) and some new ones (instant messenger, chat, forums and more), and it does so in a way that is more intuitive. Many who try it soon find that they can't live without it."
- "We use www.citadel.org. It is open source, free, rock solid and well supported. Now we have 50 accounts (but 40GB of indexed mail)."

Scalix (along with some of the others in some cases) got mixed reviews:

- "I used Scalix at the last company I worked for and liked it a lot."
- "Scalix is not open source, and the Community version is seriously broken (no CalDAV, etc.)."
- "OX/Scalix/Zimbra are bloated and pretend to be open source."

Others: Kolab, Dovecot, Gordano, Bynari, Postfix+dovecot+chandler-server, Axigen, Communicate, Mailtrust, Novell Groupwise, OpenExchange, Zafara, Sun suites. There was no mention of Scalable OpenGroupware.org, which Francis Lachapelle and Ludovic Marcotte reviewed in April 2008.

So the challenge remains. Any bets on how long it will take to finish the job? ■

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.

MAXIMIZE YOUR RETURN ON INVESTMENT AND INFORMATION BY SQUEEZING ALL YOU CAN OUT OF EVERY GIGABYTE, TERABYTE AND PENNY.

M_H

MANAGED HOSTING

Starting at \$299/month

Fully customized and dedicated server configurations

- Customization of configurations
- Full OS environment control
- Server, storage, security and network devices
- Dedicated Support Team built around you
- Managed by certified experts 24x7x365

C_H

CLOUD HOSTING

Starting at \$20/month

Pay-as-you-grow solutions for sites, apps and email

- Load-and-go websites
- On-demand virtual Linux servers
- Fast, inexpensive web storage
- Easy deployment, scaling and online administration
- Round-the-clock management of infrastructure

E_H

EMAIL HOSTING

Starting at \$3/mailbox per month

Business-class email for companies large and small

- Hosted Microsoft® Exchange
- POP/IMAP email with Outlook capabilities
- Mobile access
- Wireless device synchronization
- Managed by email specialists at all times

Managed, Cloud and Email Hosting solutions – whatever your need and budget, the world's leader in hosting has the answer you want. Not to mention the nonstop Fanatical Support® you deserve.



Microsoft

DELL

CISCO



AMD

vmware



rackspace.com/linuxjournal • 888-571-8976 | *experience fanatical support®*

rackspace
HOSTING



GPU Computing

More GFLOPS, Fewer Watts!

Clusters with Tesla S1070 - 4 GPU Servers

- ▶ 36 GPUs + 36 CPUs + 24 TB in 24U
- ▶ 40 Gbps FasTree™ InfiniBand
- ▶ InfiniScope™ Network Monitoring

WhisperStation™ with 1 to 4 Tesla GPUs

Tesla C1060 GPU Performance:

- ▶ 1 TFLOPS per GPU
- ▶ 4 GB DDR3 per GPU
- ▶ 102 GB/Sec Bandwidth



FireStream™ Workstations & Servers

- ◉ WhisperStation™ - FS with FireStream 9250 Stream Processors



AMD FireStream 9250

Stream Processor Performance

- ▶ 1 TFLOPS
- ▶ 1 GB DDR3
- ▶ Single Wide Card /150 Watts
- ▶ Double Precision Floating Point Hardware

Stream SDK

- ▶ Brook+ Open Source Compiler
- ▶ ACML and CAL Math Libraries