

LINUX JOURNAL™

Since 1994: The Original Magazine of the Linux Community

Interview with
Christof Wittig »
& Jerry Fiddler
of **db4objects**



WINDOWS MIGRATION & INTEGRATION

So Easy...

- » MANAGING NTFS FROM LINUX
- » WIN4LIN VIRTUAL SERVER DESKTOPS
- » SWAP EXCHANGE WITH CITADEL
- » WINDOWS ACCESS TO LINUX FILESYSTEM

FEBRUARY 2007 | ISSUE 154
www.linuxjournal.com



USA \$5.00
CAN \$6.50



0 71486 03102 4



UDEV FOR
FLASH BACKUP

C++, STLDB4
AND BERKELEY DB

LIBFERRIS AND
VIRTUAL DOCUMENTS

Enterprise and High-Performance Computing Under Your Control



Industry Leading 2P and 4P x86 Computing

Innovative server technology with outstanding performance and memory scalability

2-way XtremeServer™ up to 64GB of memory



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

4-way XtremeServer™ up to 128GB of memory



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

Keeping Your Cool in the Data Center while Consolidating and Virtualizing your IT Infrastructure

Go to [http://www.appro.com/whitepaper/White Papers.asp](http://www.appro.com/whitepaper/White%20Papers.asp)

- AMD Opteron™ Processors
- Quad-Core Ready - increase capacity without altering datacenter infrastructure
 - Best performance per-watt with energy-efficient DDR2
 - Optimized system performance with Direct Connect Architecture

For more information, please visit www.appro.com
or call Appro Sales at 800-927-5464 or 408-941-8100.



Today, Dan configured a switch in London, rebooted servers in Sydney, and watched his team score the winning goal in St. Louis.

With Avocent data center solutions, the world can finally revolve around you. Avocent puts secure access and control right at your finger tips – from multi-platform servers to network routers, your local data center to branch offices, across the hall or around the globe. Let others roll crash carts to troubleshoot – with Avocent, trouble is on ice.

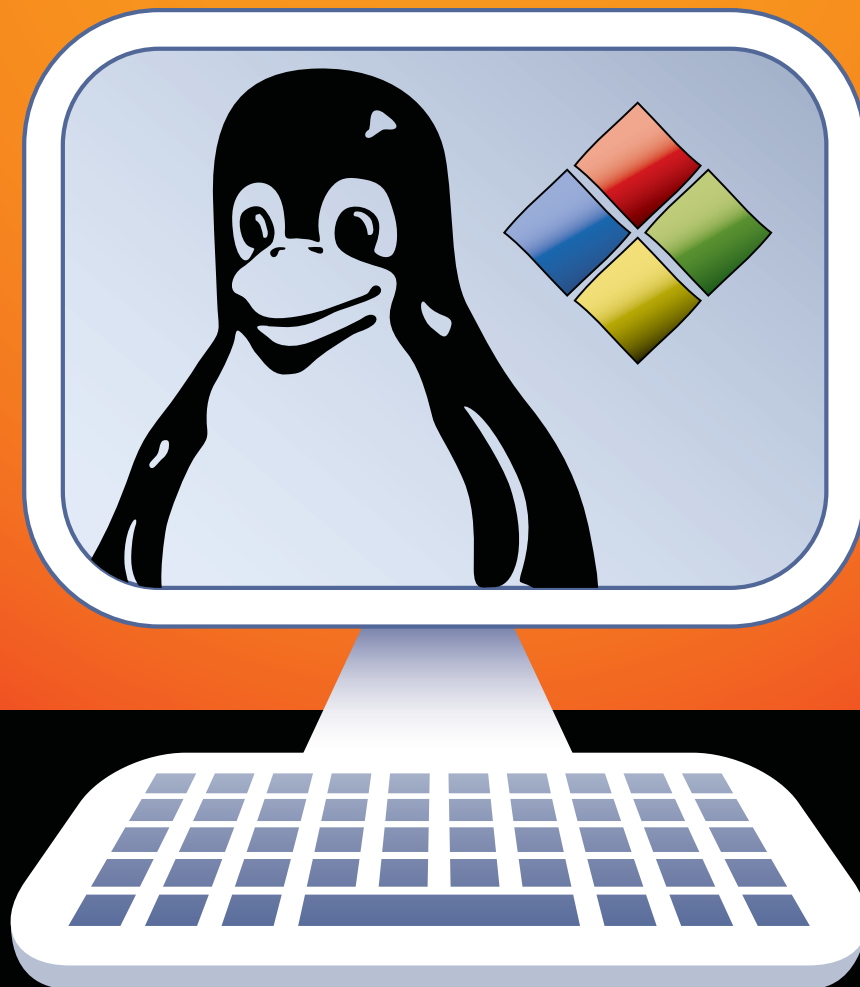
To learn more, visit us at www.avocent.com/ice to download Data Center Control: Guidelines to Achieve Centralized Management whitepaper or call 866.277.1924 for a demo today.



CONTENTS

FEBRUARY 2007

Issue 154



Migration

FEATURES

48 THE INTEROPERABILITY POWER OF LINUX-NTFS TOOLS

Linux does NTFS better than you think.

Steven Mathes

52 BREAK THE HARDWARE UPGRADE CYCLE WITH WIN4LIN WINDOWS VIRTUAL SERVER

Windows on many Linux clients.

Jon Watson

58 ACCESSING LINUX FILESYSTEMS IN WINDOWS

How do you reach Linux filesystems from Windows?

Irfan Habib

62 THE CITADEL GROUPWARE SERVER

Replace Microsoft Exchange.

Jon Watson

ON THE COVER

- Christof Wittig & Jeffy Fiddler of db4objects, p. 68
- Managing NTFS from Linux, p. 48
- Win4Lin Virtual Server Desktops, p. 52
- Swap Exchange with Citadel, p. 62
- Windows Access to Linux Filesystem, p. 58
- UDEV for Flash Backup, p. 92
- C++, STLDB4 and Berkeley DB, p. 80
- Libferris and Virtual Documents, p. 72



One

PGI Unified Binary™

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



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

CONTENTS FEBRUARY 2007

Issue 154

COLUMNS

20 TECH TIPS WITH
GNULL AND VOYD
Chester Gnull and Laverta Voyd

24 REUVEN M. LERNER'S
AT THE FORGE
Scriptaculous

28 MARCEL GAGNÉ'S
COOKING WITH LINUX
Words, Words, Words...



32 DAVE TAYLOR'S
WORK THE SHELL
Analysing Your Search Keywords

36 MICK BAUER'S
PARANOID PENGUIN
Introduction to SELinux

40 JON "MADDOG" HALL'S
BEACHHEAD
Beneath the Surface

42 DOC SEARLS'
LINUX FOR SUITS
Migrating a Mentality

96 NICHOLAS PETRELEY'S
/VAR/OPINION
Long Live the Freedom of Linux

IN EVERY ISSUE


8 LETTERS
12 UPFRONT
46 NEW PRODUCTS
81 ADVERTISERS INDEX

INDEPTH

68 INTERVIEW WITH CHRISTOF
WITTIG AND JERRY FIDDLER OF
DB4OBJECTS

Company and community overlap.

Nicholas Petreley



72 VIRTUAL FILESYSTEMS ARE
VIRTUAL OFFICE DOCUMENTS

Are libferris features virtually inexhaustable?

Ben Martin

80 SIMPLE ACCESS BERKELEY DB
USING STLDB4

Tired of the thin C++ wrapper for Berkeley DB?

Ben Martin

86 CREATING SELINUX POLICIES
SIMPLIFIED

SELinux is easier than you think.

Irfan Habib

89 INTEGRATING PHP AND PERL
PHP or Perl? You don't have to decide.

Irfan Habib

92 PAINLESS THUMBDRIVE
BACKUPS

UDEV makes Flash backups easy

Andrew Fabbro



12 UPFRONT: A WHITE BOX PHONE

Next Month

VOIP

Put an asterisk next to next month's issue, because we have a variety of VoIP tutorials for using Asterisk in your PC or even an embedded device. We'll also tell you how to set up Trixbox for your VoIP system and show you how to get CommuniGate Pro VoIP working.

As always, there's much more. We'll give you a primer on the D language, an extremely powerful compiled yet managed language, garbage collector and all. You'll also find out how to build a multiroom digital music system controlled by a Nokia 770.

USPS LINUX JOURNAL (ISSN 1075-3583) is published monthly by Belltown Media, Inc., 2211 Norfolk, Ste 514, Houston, TX 77098 USA. Periodicals postage paid at Houston, Texas and at additional mailing offices. Cover price is \$5 US. Subscription rate is \$25/year in the United States, \$32 in Canada and Mexico, \$62 elsewhere. POSTMASTER: Please send address changes to *Linux Journal*, PO Box 980985, Houston, TX 77098. Subscriptions start with the next issue.

Gemini

2U

– Ultra Dense Series



Up to
**40% Less Power
Consumption**

Up to
40% Less Heat

- Two fully independent systems in a 2U
- Ability to run two discrete operating systems in one box
- Up to 8 CPU cores (forward compatible to 16)
- Up to 12 hot-swappable SATA, SCSI, or SAS hard drives
- RAID 0, 1, 5, 6, 10, 50 available on both systems
- Opteron™ or Xeon™ multi-core processors
- Up to 64GB memory per motherboard
- One available PCI-X or PCI-E slot per motherboard



Patent Pending

2 Systems in One

Built on open standards, the Gemini 2U elegantly accommodates two discrete motherboards in a 25" chassis uniquely designed for easy access from the rear.

Gemini 2U represents the realization of intoxicating power and superior environmental specifications, with up to 40% less power consumption, 40% less heat and 30% less noise. Remarkably, it all fits nicely into any standard rack. Front to back, Gemini 2U is both powerful and efficient.

At Open Source Systems we understand you need practical, customizable, and affordable solutions that are easy to manage and maintain.

For more information and to request your evaluation unit today, visit us at www.OpenSourceSystems.com, or call direct at 866.664.7867.



 **Open Source
Systems™**

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

Editor in Chief

Nick Petreley, ljeditor@linuxjournal.com

Executive Editor

Jill Franklin
jill@linuxjournal.com

Senior Editor

Doc Searls
doc@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

Contributing Editors

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

Proofreader

Geri Gale

Publisher

Carlie Fairchild
publisher@linuxjournal.com

General Manager

Rebecca Cassity
rebecca@linuxjournal.com

Director of Sales

Laura Whiteman
laura@linuxjournal.com

Regional Sales Manager

Joseph Krack
joseph@linuxjournal.com

Circulation Director

Mark Irgang
mark@linuxjournal.com

Marketing Coordinator

Lana Newlander
mktg@linuxjournal.com

System Administrator

Mitch Frazier
sysadm@linuxjournal.com

Webmaster

Keith Daniels
webmaster@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

Editorial Advisory Board

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

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: +1 713-589-3503
FAX: +1 713-589-2677
TOLL-FREE: 1-888-66-LINUX
MAIL: PO Box 980985, Houston, TX 77098 USA
Please allow 4-6 weeks for processing address changes and orders
PRINTED IN USA

LINUX is a registered trademark of Linus Torvalds.



TYANPSC™

Hands On Supercomputing



Typhoon™ 600 Series Personal Supercomputer

TyanPSC's Typhoon™, the next generation turnkey Personal Supercomputer has the power to blow away all your computational needs! Purpose-built for office and laboratory environments, easy to deploy and use, the Typhoon™ provides intense computational power in remote or constrained places, works like a PC and is whisper quiet.

High Performance Computing Just Got Cooler

Clusters of Typhoons / Low Power
Small size Form Factor / Under Mobility



Typhoon

T-630 DX / T-650 QX Series

- Up to 186 / 256 GFlops at your desk!
- Turnkey, Easy-to-Deploy, and Easy-to-Use
- Integrated 5 node cluster - up to 20 / 40 processor cores in a box!
- Plugs into standard wall outlet - only uses 15 Amps
- Microsoft® Windows® Compute Cluster Server 2003 pre-installed
- High Performance in Constrained spaces: Office, Remote, Plane, Boat, etc
- RAID capable



**Windows Compute
Cluster Server 2003**

TYANPSC™
Personal Supercomputer

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

For More Information, please visit
www.tyanpsc.com

Letters



Kudos for the Bold Call to Boycott Novell

I wanted to write since it seems you're getting some flak in response to your well articulated blog. You probably already know this well enough, but it is something that Vaughan-Nichols has been publishing, and I have been trying to convey in similar published articles. Only those who've observed MS business strategy over the years seem to appreciate the key elements of your statement.

So, to help counteract some of that flak, I wanted to say THANK YOU for holding a position, standing on substance, and being candid.

--
Mark Rais

Note: This letter is in response to blog articles published exclusively on the www.linuxjournal.com site: "Novell Is Loading Microsoft's Gun", www.linuxjournal.com/node/1000129 and "A Five Year Deal with Microsoft to Dump Novell/SUSE", www.linuxjournal.com/node/1000121.—Ed.

Marcel Gagné's Column

In the November 2006 Letters section, Marcel Gagné writes, "With a very few exceptions (such as Mr DeSouza), I get nothing but praise for my articles."

Let me cause another exception for Mr Gagné and add my voice to those who think the style of Mr Gagné's articles are extremely annoying. Usually I just scan the article looking for the actual content, as the fluff is too cheesy and bothersome to wade through. Asking around my department (there are many avid readers of *Linux Journal* here), no one else even bothers to read the article because of the style.

--
Chris Russell

With apologies to Marcel for inadequately borrowing his style, François serves a sweet white wine to some and a rich full-bodied red to others. Most of our readers love Marcel's column. Your exception is respected and noted.—Ed.

Letters to the Editor

Every time I was starting a new project, *Linux Journal* beat me to it with an introduction article. When I wanted to go PPC, my new issue just arrived. When I was reading about Qtopia, my new issue arrived. When they introduced the Nokia 770, my issue arrived to tell me all about its UI.

I don't expect every article to be meaningful for me, but I do expect something from every issue. Telling us about x264 was great (the part about QuickTime users in specific), but the rest of the issue needs to be refocused.

What I would like to see? How about a technical discussion of how the Zaurus "Sharp ROM" is put together? How about a discussion on how to get the 770's window system working for a workstation? How about an article about the Qtopia phone? Does Trolltech offer a GPL phone edition for use with the new Wi-Fi IP SIP phones (for example, WIP300)? I could go on and on about the growing trends in these directions.

--
J.

Thanks for the suggestions. We welcome input like this.—Ed.

P2P over NAT

I enjoyed Girish Venkatachalam's article in the August 2006 issue of *LJ* about developing P2P protocols across NAT. I was particularly interested in some assertions made and was wondering if Girish can provide some references to those.

In particular, "At least 50% of the most commonly used networking applications use peer-to-peer technology." This doesn't seem right.

I always thought that, being connection-based, TCP was a lot easier to NAT than UDP, although some TCP applications make it harder by including IP addresses in the

Optimal awking

I admit that I did not read the original article "Analyzing Log Files" [October 2006 issue of *LJ*] by Dave Taylor. I did see the "Optimal awking" letter however [Letters, December 2006 issue of *LJ*]. Being an old "bit twiddler", I was interested in the enhanced run-time mentioned by reordering the original:

```
awk '{ print }' access_log | sort | uniq -c | \
  sort -rn | grep "\.html" | head
```

to:

```
awk '{ print }' access_log | grep '\.html' | sort \
  | uniq -c | sort -rn | head
```

Now, I'm not really an awk person, but I was curious as to what the awk program did. Apparently, it is just an expensive version of cat—that is, it copies

its input to stdout, unchanged. In that case, why even have it? Also, why use grep? Instead, use fgrep, which, in this case, produces the same result with somewhat less overhead:

```
fgrep '.html' access_log | sort | uniq -c | sort -rn | head
```

[This] should produce the same output and totally eliminate the awk. In this case, no big deal. Unfortunately, many neophytes will pick up a script from a magazine and use it without really understanding it. So, I am a bit picky about examples. For a one-shot, this is no big deal. But I am a bug for efficiency—comes from programming back on 1MHz 8080s, I guess.

Unfortunately, I don't have a Web access_log to do any testing to see if this really makes much of a difference.

--
John McKown

Linux Laptops

Starting at \$799



Linux Desktops

Starting at \$375



Linux Servers

Starting at \$899



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



awk Is Not So Awkward

I don't think Dave Taylor gives awk enough credit [see Dave's November 2006 column]. I do not have access to the same Web files, logs or version of Linux. However, I do know that his solution can be written entirely with awk. Using AIX and HP-UX, I did dummy up a mail log file, cheated on the date command and tested my awk solution.

Below is awk code that I think would duplicate Dave's example:

```
#!/bin/sh

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

awk '
    ( index($0, YESTERDAY) ) {
        hits++;
        bytes+=$10
        next}
END {
    printf("Calculating transfer data for %s\n", YESTERDAY)
    printf("Sent %d bytes of data across %d hits \n", bytes, hits)
    printf("For an average of %d bytes/hit\n", (bytes / hits) )
    printf("Estimated monthly transfer rate: %d \n", (bytes * 30) )
}
' YESTERDAY="$(date -v-1d +%d/%b/%Y)" ${LOGFILE}
--
Rory Hammond
```

application layer part of the datagrams (such as FTP in "normal" mode), requiring the NAT router to have to inspect and modify every FTP datagram.

--
Bob

SSH and DenyHosts

Regarding the yellow warning about SSH dictionary attacks [LJ, December 2006, "A Server (Almost) of Your Own" by George Belotsky], try the wonderful DenyHosts (www.denyhosts.net).

DenyHosts monitors the incoming connections into your server (mainly SSH, but it can be FTP, POP or anything else with a login/password and log file) and blocks source IP addresses by automatically adding entries to /etc/hosts.deny.

I put it to run on a brand new Web server a month ago, and it has already more than 7,000 forbidden addresses!

Besides, you also can share your blacklist with DenyHosts' Web site, feeding a mega-blacklist of the really bad guys.

--
Carlos Vidal

Java Java Java

I've been waiting and waiting and waiting for Java articles. How about Java applets (Ajax before Ajax), Java servlets, Swing (and real cross-platform stuff). I look forward to hearing and reading about it.

--
Mark Molenda

We have just such an article in the works!—Ed.

PlayStation Games on Linux

In the May 2006 issue of LJ, there was an article by Dee-Ann LeBlanc on the above subject.

Unfortunately, the emulator that Dee-Ann recommended is still available but unsupported for Linux; their Linux guy left the project. *But*, all is not lost; Linux will not be beaten. There is a new Linux emulator starting called PCSX2, and it can be found at pcsxii4unix.sourceforge.net. The new version is not completed yet; they need help.

Thanks LJ for the best Linux mag on the continent. Always remember: there are Linux users, and then there is the rest of society.

--
Des Cavin

One Linux, Many Faces

In his letter [December 2006 issue of LJ], Jon Alexander described how he surprised his friend by logging in with different desktops. If he really wanted to impress his friend, he could have logged in with multiple desktops and then switched between them. He even could have included a couple of remote desktops for good measure. Also, don't forget about the virtual desktops that Linux supports.

--
James Knott

Thieving Teachers

It was interesting to read Jon "maddog" Hall's experience in his article titled "Soweto: Power from the People" [see the UpFront section of the December 2006 issue of LJ] and compare it with an experience our computer club had a few years ago.

The club committee decided that if the club members would donate their redundant computer equipment to the club, the club would donate this equipment to a Soweto school and help them set it up as a local network and subsequently connect to the Internet.

The installation went fine, and the local network

MAGAZINE

PRINT SUBSCRIPTIONS: Renewing your subscription, changing your address, paying your invoice, viewing your account details or other subscription inquiries can instantly be done on-line, www.linuxjournal.com/subs. Alternatively, within the U.S. and Canada, you may call us toll-free 1-888-66-LINUX (54689), or internationally +1-713-589-2677. E-mail us at subs@linuxjournal.com or reach us via postal mail, Linux Journal, PO Box 980985, Houston, TX 77098-0985 USA. Please remember to include your complete name and address when contacting us.

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

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

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

ADVERTISING: *Linux Journal* is a great resource for readers and advertisers alike. Request a media kit, view our current editorial calendar and advertising due dates, or learn more about other advertising and marketing opportunities by visiting us on-line, www.linuxjournal.com/advertising. Contact us directly for further information, ads@linuxjournal.com or +1 713-344-1956 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.

A Touch of Elegance

I hesitate to take issue with an authority like Dave Taylor, but [his script] uses repeated (and in many cases redundant) system calls and divisions to achieve what simple multiplications can do [see Dave's December 2006 column]. It also converts the results to the wrong units. (See `man units` for a discussion.)

A simple algorithm to achieve the same effect is embedded in a test harness as follows:

```
#!/bin/sh
# Script for numeric scaling - $1 = number, $2 = iterations

for (( i = 1; i <= $2; i++ ))
do
    ki=1024
    mi=$(( $ki*$ki ))      # 1048576    without typo risks
    gi=$(( $ki*$mi ))     # 1073741824 without typos

    value=$1
    if [ $value -lt $ki ] ; then
        units="bytes"
    elif [ $value -lt $mi ] ; then
        units="KiB"
        div=$ki          # < 1 Mi, so calculate Kibytes
    elif [ $value -lt $gi ] ; then
        units="MiB"
        div=$mi          # < 1 gi, so calculate Megs
    else
        units="GiB"
        div=$gi          # >=1 gi, so calculate Gigs
    fi

    if [ $units != "bytes" ] ; then # scale value appropriately
        value=$(echo "scale = 2; $value / $div" | bc )
    fi

    echo "$value $units"
done

#      End tcon2
```

Running 1,000 iterations of each on an HP laptop with an AMD 2500 chip showed the revised version to take approximately one-quarter of the time (real, user, and system) of the original.

--

Alan Rocker

also worked; however, when the installation team arrived to connect the school to the Internet, what did they find? Every single piece of equipment had either been stolen or broken. It was subsequently discovered that some of the teachers were responsible for some of the theft.

--

Alf Stockton

Why Ubuntu and Then KDE?

I found it a little bit ironic that you chose Ubuntu as the 2006 Editors' Choice Linux distribution, but that you chose KDE as your Editors' Choice desktop environment. It seems to me that if you were going to pick Ubuntu, you'd choose GNOME, and if you were going to choose KDE, you would choose Kubuntu. Are there reasons you picked Ubuntu over Kubuntu, or did you simply mean (K)ubuntu in general for your Editors' Choice distribution?

PS. Long live KDE!

--

Geoff

As we said in our write-up, we also find it a puzzler as to why Ubuntu seems to be a favorite, yet research data shows people prefer KDE over GNOME by a significant margin. Perhaps people refer to all variants of Ubuntu as Ubuntu, even if what they're really using is Kubuntu. Or, maybe others do like some of us at Linux Journal do. Some of us at Linux Journal install Ubuntu and then install and use KDE (thus essentially converting it to Kubuntu).—Ed.

Erratum

Some of the code was inadvertently formatted incorrectly in George Belotsky's "A Server (Almost) of Your Own" in our December 2006 issue. For the corrected version of the article, please see www.linuxjournal.com/article/8337.

VERIO HOSTING IS LINUX WITH A LINEAGE.

- **Root Access:** Providing the control you need.
- **Advanced FairShare Technology:** Better resource management means better performance.
- **Support That's Actually Supportive:** Award-winning support provided by system administrators.

Announcing Verio Linux[®] VPS.

At Verio, we have a long-standing commitment to open source, dating back to our roots in FreeBSD. Now, as the pioneer in virtual private server (VPS) technology and as a hosting provider backed by the financial resources of the world's largest telecommunications company, we bring something extra to Linux: reliability. To learn more, call 1-877-837-4654 or visit www.verio.com/linuxlineage.

There is no substitute for the right foundation.

Verio and the Verio logo are trademarks and/or service marks of Verio Inc. in the United States and other countries. Linux[®] is the registered trademark of Linus Torvalds in the U.S. and other countries. All other names are trademarks or registered marks of their respective owners. ©2006 Verio Inc. All rights reserved.

Build on us.

VERIO
An NTT Communications Company

diff -u

WHAT'S NEW IN KERNEL DEVELOPMENT

"I'm not a huge fan of the LGPL, especially with the recent issues of GPLv3. The reason? The LGPL is expressly designed to be compatible with the

GPL, but it's designed to be compatible with any version (and you can't limit it, the way you can the real GPL). So you can take LGPL 2.1 code, and relicense it under GPLv3, and make changes to it, and those changes won't be available to a GPLv2 project."
—Linus Torvalds

The **sysctl** call, allowing users to configure kernel parameters at runtime, is likely to go away. This goes against the standard doctrine of never breaking user space, but the kernel folks may get away with it this time, because it looks as though there are no user-space programs that actually use sysctl. Apparently, people do their kernel configuration operations in other ways. If you or someone you love depends on sysctl, you might consider raising the issue on the linux-kernel mailing list while there's still time. **Linus Torvalds** and **Andrew Morton** have both expressed the opinion that taking sysctl out would be the right thing to do—Linus because no one uses it and Andrew because it would be a shame to leave a big wad of such useless code in the kernel permanently, if a viable alternative existed. But, in case it really would just break too much stuff, **Albert Cahalan** has volunteered to be the official sysctl maintainer if one is needed.

The Multimedia Card subsystem is now the **Multimedia Card and Secure Digital subsystem**, and **Pierre Ossman** has submitted a patch making himself the new maintainer. **Russell King**, the previous maintainer, had stepped down and marked the subsystem "orphaned". Meanwhile, **Jiri Slaby** has added new maintainer entries for the **Moxa SmartIO/IndustIO Serial Card driver** and the

Multitech Multiport Card driver, in both cases naming himself as the official maintainer. An anonymous kernel tester has reported some benchmarks showing that **ext4** is about 20% faster than either **ext3** or **Reiser4**. Although a useful (and perhaps gratifying) result, **Theodore Ts'o** pointed out that what was really needed was an automated testing infrastructure, so that each version of each filesystem could be compared, and the particular results correlated to the specific patch that either sped things up or slowed things down. And, various other folks suggested incorporating tests for other filesystems as well. The original poster agreed that this would be great, but he or she (and Ted) also pointed out that the amount of work required to create such an infrastructure would be massively big. It does not look as though an automated filesystem benchmark is coming any time soon, though you never know.

It's apparently **wiki season** in kernel land. **Valerie Henson** has created two wikis, one for filesystems at **linuxfs.pbwiki.com** and the other for huge memory pages at **linux-mm.org/HugePages**. As one might expect, the filesystem wiki is a bit more active than the huge pages wiki. To go along with these collaborative projects, Valerie has also started up two **IRC channels** on **irc.oftc.net**: **#linuxfs** and **#hugepages**. Meanwhile, **Darren Hart** and **Theodore Ts'o** have started up a wiki for real-time support at **rt.wiki.kernel.org**, and in fact, the generic **wiki.kernel.org** site may offer generic wiki hosting services to any legitimate kernel project. Just ask the site administrators to set it up for you! At the same time, as Ted points out, you should make sure there is at least a person or two to act as editor and maintainer, or your wiki is likely to become stale. Nothing like stale wiki to clear the sinuses, I say!

—ZACK BROWN

LJ Index, February 2007

1. Number of Gannett newspapers that will "crowdsource" editorial from readers and bloggers: **91**
2. Number of US intelligence agencies that will join to create "Intellipedia": **16**
3. Number of ways "distributive networks" of citizen journalists covered the November 2006 elections: **9**
4. Percentage of US businesses with fewer than ten employees that don't have better than dial-up Net access: **60**
5. Minimum thousands of certificates Microsoft will distribute per year allowing its customers to use SUSE Linux: **70**
6. Minimum millions of dollars Novell will receive from Microsoft for those certificates: **240**
7. Rounded thousands in dollars per certificate: **3.4**
8. Linux-based hosters among Netcraft's top 50 most reliable hosting providers for November 16, 2006: **23**
9. Windows-based hosters among Netcraft's top 50 most reliable hosting providers for November 16, 2006: **11**
10. BSD-based hosters among Netcraft's top 50 most reliable hosting providers for November 16, 2006: **7**
11. Solaris-based hosters among Netcraft's top 50 most reliable hosting providers for November 16, 2006: **4**
12. Number of Linux-based hosters among Netcraft's top four for October 2006: **4**
13. Number of Linux-based hosters among Netcraft's top ten for October 2006: **5**
14. Number of FreeBSD-based hosters among Netcraft's top ten for October 2006: **2**
15. Number of Solaris-based hosters among Netcraft's top ten for October 2006: **3**
16. Number of Windows-based hosters among Netcraft's top ten for October 2006: **0**
17. Position of IBM's Linux-based Blue Gene/L (at Lawrence Livermore National Laboratory) among Top 500 Supercomputers for 2006: **1**
18. Trillions of calculations per second measured on Blue Gene/L: **280.6**
19. Other new Linux-based supercomputer clusters at Lawrence Livermore National Laboratory: **4**
20. Price in millions for Lawrence Livermore National Laboratory's four new supercomputer clusters: **11**

Sources: 1: WIRED NEWS | 2: WASHINGTON POST | 3: NEWASSIGNMENT.NET
4: OECD.ORG (JUNE 2006 STATS) | 5-7: SILICON REPUBLIC.COM
8-11: NETCRAFT.COM (DURING THE PAST 24 HOURS FROM THE DATE LISTED)
12-16: NETCRAFT.COM | 17, 18: TOP500.COM AND IBM | 19, 20: INFORMATIONWEEK
—Doc Searls

The Industry Leader for Server Appliances



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

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

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

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



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

A White Box Phone

As we know too well, embedding Linux in a device doesn't make it "open". And although there are open Linux-based embedded devices, telephones fitting that description have been rare.

The OpenMoko phone aims to change that (openmoko.com). Funambol (pronounced *fo-nahm-ball*), a Taiwan-based manufacturer whose official ambition is "to bring the customer benefits of open source software to the \$300 billion global mobile market", launched OpenMoko to generally positive reviews. From my own contact list, these ranged from Gordon Cook's (gordoncook.net) "This is AMAZING STUFF" (in fact, I heard about it first from Gordon) to Bob Frankston's (frankston.com) "No Wi-Fi? Huh?" and "As I've discovered with my current programmable phone, having Wi-Fi and GPS can make a big difference."

But the quote that matters most comes from Harald Welte, who wrote this in his blog at gnumonks.org (gnumonks.org/~laforge/weblog/2006/11/08/#20061108-my_no_longer_secret_project):

In this project I'm responsible for the system-level software design and implementation. This means: kernel, drivers, GSM communication infrastructure, etc.

So why is this project so exciting? Because it's [yet another] Linux phone? No. It's because this is the first time (to the best of my knowledge), that a vendor is:

- » involving (hiring) prominent community members to do the actual architecture design and implementation;



- » planning to completely open up their Linux distribution for any contributed development, e.g., use a package manager that can access arbitrary package feeds;
- » trying very hard to make sure almost everything will be Free Software, from drivers up to the UI applications;
- » actively providing documentation and interfaces for third-party development on any level of the system, from debug interface, boot loader, kernel, middleware through the UI applications;
- » using X11 to allow users to run any existing X11 Linux application (within resource constraints).

So basically, from a Free Software community level, this is exactly the kind of phone you want to get involved with, and play with. Yes, it's not the perfect phone. It runs a proprietary GSM stack on a separate processor. There are some minor, self-contained proprietary bits on the back end side in userspace. But well, it's probably the best you can do as a first shot of a new generation of devices, and without too much existing market power to put on upstream vendors.

Beats anything you'll ever read in a press release. I'll give the last word to Brad Fitzpatrick (brad.livejournal.com), father of LiveJournal OpenID, memcached and other fine hacks. After news of OpenMoko hit the streets, Brad wrote, "On sale Jan 2007...I'm totally getting one".

—DOC SEARLS

They Said It

Companies come and go but you only get one reputation.
—David Sifry (to Doc Searls on the phone)

Good Web 2.0 sites follow the UNIX design model: do one thing well, and play well with others.
—Evan Prodromou, www.linuxworld.com/news/2006/110906-web20-openid.html

Coding up the simplest thing that could possibly work is really about this: If you can't keep five things in your head at one time and make a decision, try keeping three things in your head. Try keeping just one thing in your head, and see if you can make a decision. Then you can think of the next thing. And amazingly, when you write some of this dumb, straight-ahead code, it often turns out that it was all that was required. It works great. When a second programmer comes back later and reads the code she might say, "The people who wrote this are morons. They just wrote a simple linear search here. This thing's ordered, so they could have done a binary search. They could have used a hash table here. Why are they doing a linear search?" Well, because a linear search worked. And when the other programmer looked at the linear search, she understood it in a minute.
—Ward Cunningham, www.artima.com/intv/simplest3.html

Men occasionally stumble over the truth, but most of them pick themselves up and hurry off as if nothing ever happened.
—Winston Churchill, www.brainyquote.com/quotes/quotes/w/winstonchu135270.html

LAMP Gets a J

You can't add the letter J to LAMP and spell anything sensible. So, some took to calling Linux "stacks" with Java "LAMJ", "LAMPJ" or "LAMP-J". But they never seemed legitimate—not so long as Java didn't have an open-source license that passed muster with the rest of the L+ alphabet.

That changed on November 13, 2006 (as we go to press here), when Sun finally announced that it would be releasing Java under the GPL—specifically under version 2, which is the license Linux has used since it came out and has stuck with (at Linus' insistence), even after version 3 was announced last year.

Sun has hinted for some time that it would go with the GPL for Java. Jonathan Schwartz, the company's CEO, hinted as much in a conversation I had

with him on stage at the Syndicate conference in December 2005. Now Jim Thompson has another intriguing question: "Is Solaris going to be the original GPLv3 *nix platform?"

In his blog, Jonathan writes, "The GPL is the same license used to manage the evolution of GNU/Linux—in choosing the GPL, we've opened the door to comingling the communities, and the code itself. (And yes, we picked GPL version 2—version 3 isn't available, but we like where the FSF is headed.)"

Whether or not Solaris and the FSF arrive at the same place, Jim has one more question: "What are the chances that Ubuntu would offer a version of its distro with a GPL-ed Solaris kernel underneath?"

Redraw your own conclusions.



Illustration by appleuz

- » Jonathan Schwartz' blog: blogs.sun.com/jonathan
- » Jonathan Schwartz' blog post: blogs.sun.com/jonathan/entry/fueling_the_network_effect
- » Jim Thompson: www.smallworks.com

—DOC SEARLS

Linux laptops. Supported.

Since 1999, EmperorLinux has provided pre-installed Linux laptop solutions to universities, corporations, and individual Linux enthusiasts. We specialize in the installation of the Linux operating system on a wide range of the finest laptops and notebooks made by IBM, Lenovo, Dell, Sharp, Sony, and Panasonic. We offer a range of the latest Linux distributions, as well as Windows dual boot options. We customize each Linux distribution to the particular machine it will run upon and provide support for: ethernet, wireless, EVDO mobile broadband, PCMCIA, USB, FireWire, X-server and 3D, sound, power management, DVD±RW, and more. All our systems come with one year of Linux technical support by both phone and e-mail, and full manufacturers' warranties apply.

Visit www.EmperorLinux.com or call 1-888-651-6686 for details.



EmperorLinux

...where Linux & laptops converge
www.EmperorLinux.com/1-888-651-6686

Model specifications and availability may vary.



Learn about specific models later in this issue

The Almost Inevitable Migration for All

This *Linux Journal* issue is about migration, but migration means many things to many people. Few companies and even few homes use only one operating system.

A lot of Linux advocates use Apple notebooks running OS X, a BSD derivative. Many of these Linux advocates use iPods or BlackBerrys. Many are forced to use Windows systems due to work pressures.

Likewise, a certain company in Redmond, Washington, often talks about the “cost” of migrating to Linux from its own operating system, and uses this “cost” to boost the Total Cost of Ownership (TCO) of our favorite operating system and design strategy of Free Software, without acknowledging the higher value of end users having control over their own destinies.

This same company also ignores the fact that it has forced more “migrations” over the years by making its customers move from DOS to Windows 3.1, Windows 95, Windows 98, Windows NT, Windows XP, Windows 2000 and Win ME and are now looking for an even greater migration to the 64-bit Vista. Although we can hope that this company has learned from all of the other 64-bit operating systems and their migration issues (Linux moved to 64-bit in 1995), we can assume that there will be some hiccups along the way. And, this seems to be borne out in the stages and delays that have come along with Vista (probably one of the most “beta-ed” of all operating systems), with cautions to companies from the producers of Vista to “test, test, test”.

Therefore, in reality, migration is really integration, unless you are lucky enough to be able to start from scratch with a one-operating-system strategy and maintain it throughout time. I call this the one-egg, one-basket mentality, and to do it with an operating system that you have zero control over is just plain suicide.

So, when I talk to people about migrating to Linux, I work on several levels. I tell them first to do “the easy stuff”.

Doing the Easy Stuff

The first part of moving to a Free Software strategy in your environment is to start getting used to Free Software while you analyze your needs. Note that I do not differentiate between personal use or corporate use when I say “analyze your needs”, because in reality, the procedure is the same for both. Only the size and complexity of the project may differ.

First, start to learn about Free Software. Go to your local bookstore, go to the operating system section, pick out a few books, go to the coffee shop that is attached to the bookstore and look through the selections for the one or two books that can help you get started in Free Software.

Ask around your development group or your local university to see whether there is a Linux

User Group near you, and sign up for the mailing list. Do not worry if you are a newbie to the list. Simply read the list, look at the archives, and if someone asks you a question, just look wise and say, “Yes, I will go along with that.”

While you are learning about Free Software, look into the objectives of organizations like the Free Standards Group (FSG) and the Linux Professional Institute (LPI). The FSG talks about the importance of written standards and how to ensure that you are not locked into a specific version of any distribution of Free Software, and the LPI tells you the breadth of information your system administrators will need to know to maintain your Free Software operating systems.

Other organizations to investigate are the Free Software Foundation, the Open Source Initiative and other community organizations, to get a better idea of what Free Software is all about.

Next, list your activities and needs. Do not say, “I need brand-name this and brand-name that.” Instead, list the needs as more generic things. For example, “I need a word processor, but I do not need a presentation package.” Or, “I need a database, but it does not have to be relational.” Or, “I need a database, and it needs to be object-oriented.” If you start listing your needs on a generic basis, you may find you can deal with a much simpler, lighter-weight solution than you originally imagined. You also may find that this solution fits a smaller system with less memory and CPU needs.

And, you also may find that certain parts of your organization or home have different needs than other parts do. You then can make a decision to use a more focused solution for one particular need or a more general solution for all of the other needs.

While you are focusing on determining the needs of the solutions, make sure you evaluate future growth and things like security, availability and scalability.

In addition, as you start to think about future needs, consider hiring a Free Software developer or a system administrator that is familiar with Free Software. All other things being equal, Free Software people will be easier for ensuring quality (due to the openness of the source code used in their projects, the mailing-list entries and so forth) and also will help generate community interest that might leverage your company’s solutions.

A friend of mine who was a system administrator for a large company was also a Free Software person. Every day he would write Free Software to help him do his job, and every night he would go home, sit beside his spouse on the couch, and while she watched TV, he would write additional code and submit it to the source pool. The next day, he would go in and find that a lot of other people were doing

the same thing. His comment to me was: “maddog, it is like speaking into a megaphone....I say so little and I get back so much.”

[Just to show that I am not chauvinistic in this case, I’d like to point out that there are female system administrators who go home and sit beside their husbands and code while their husbands watch TV.]

After you have determined your needs, you now can start to think about alternatives in cost.

Trade-offs in Cost

It is not without reason that some of the first uses of Free Software were in the use of generalized appliances—machines and systems that end users did not see or care what the operating systems were, as long as they were stable, scalable, secure and inexpensive. These appliances manifested themselves in DNS servers, firewalls, routers, Web servers and file-and-print servers.

Why should people put a DNS server on the same machine as their highly tuned, high-performance hardware database machine? Why not split that functionality off to a smaller, less-expensive (and perhaps older) dedicated system? Even if you are a virtualization fan, the idea of using a different partition for your DNS server allows for separation of function, which in turn may allow for greater stability between components.

People setting up Web server farms quickly learned that their customers could not tell the difference between a highly expensive proprietary system serving up Web pages from a much less expensive commodity hardware solution running Free Software, other than the fact that the price/performance was better, and therefore, allowed for more machine power and greater overall reliability.

Database companies were able to sell a total solution to their customers using a free operating system running on less expensive hardware, and end-user client programs could not tell the difference in data coming over the Internet.

Of course, there are also database solutions that are recognized to be Free Software in and of themselves. MySQL and PostgreSQL are two of them. Being careful to utilize standard interfaces and commands gives you the most freedom when using any of the database products or projects.

File-and-print servers could be set up that supported not only Windows clients, but Apple clients, UNIX clients and Linux clients with the same server at the same time, invisible to the end user.

A large health company in Australia was interviewed in 1996 about whether it used Free Software. The interviewee was told “No” by the CIO, that the company did “important things” and would not use “hobbyist software” to do them. Unfortunately for that CIO, his staff members had been told to use Windows NT for a file-and-print server, and after failing a

Hear Yourself Think Again!



WhisperStation™ **Cool... Fast... Silent!**

For 64-bit HPC, Gaming and Graphic Design Applications

Originally designed for a group of power hungry, demanding engineers in the automotive industry, WhisperStation™ incorporates two dual core AMD Opteron™ or Intel® EM64T™ processors, ultra-quiet fans and power supplies, plus internal sound-proofing that produce a powerful, but silent, computational platform. The WhisperStation™ comes standard with 2 GB high speed memory, an NVIDIA e-GeForce or Quadro PCI Express graphics adapter, and 20" LCD display. It can be configured to your exact hardware specification with any Linux distribution. RAID is also available. WhisperStation™ will also make a system administrator very happy, when used as a master node for a Microway cluster! Visit www.microway.com for more technical information.

Experience the "Sound of Silence".

Call our technical sales team at 508-746-7341 and design your personalized WhisperStation™ today.



Microway
Technology you can count on™

number of times (and experiencing his wrath), they turned to a Free Software operating system, and by that time had been using it for six months. When asked when they were going to tell their CIO that they had been using “hobbyist software” to do “important things”, they estimated that “another six months of flawless operation” would do the trick.

This explains why reports from analyst companies had a “step function” in Free Software usage in the 1998–2000 era. The analysts stopped surveying CIOs and started talking to system administrators, who actually were implementing the solutions with Free Software. When the system administrators confessed to using Free Software, the charts produced by analysts showed drastic change.

In all of these areas, Free Software is mostly or totally invisible to the real end user (including home end users), and at most it requires training of system administrators to configure and set up the systems.

The Already-Functioning Desktop

Another area where Free Software can be used is on the already-existing desktop. Again, if you list the functionality needed for the “job”, instead of brand names, often a more robust solution appears.

An obvious solution is “Web browser” rather than Internet Explorer. Various Web browsers exist in the Free Software community, and each has its own advantages. Some are smaller and easier to embed. Some use less eye candy and are easier to use on the small screen (or leave more real estate for other applications on the larger screen). Some are more portable across various operating systems, so if your people are moving from one system (OS X to Windows, for example) you may want to use a browser that works well in many environments.

Other areas of upper-level compatibility are things like word processors. For the most part, I use OpenOffice.org. It works on all the operating systems that I would have wanted to use in the past ten years: Windows, Linux (including Alpha Linux), Solaris and FreeBSD.

I often questioned why someone would want to use an office system that ran on only one operating system, or even two? I found it difficult to live with needing two operating systems on my

desk—one to do my work and one to communicate with my management and sales staff. Today, I need only one system on my desk, because my solutions run across multiple operating systems.

Many Free Software solutions run on multiple operating systems. The GNU compilers, for instance, have been providing programmers with an excellent set of tools for more than 20 years. They have allowed programmers to concentrate on the basic algorithm without having to worry about the incompatibilities of syntax and semantics that can occur across compilers written by different organizations and for different hardware architectures. It is true that some commercial companies do the same things with their very excellent commercial compilers. This provides the end user with customer choice.

All Pain, No Gain

So far, I have been talking about everything but the cold-turkey movement from a proprietary solution to a Free Software solution. Now I am going to say something that will (I am sure) surprise a lot of Free Software people.

If you have a solution that is working fine for you, is incredibly stable, has no bugs, is reasonable in price, is from a solid company that is not looking to change its products radically (thus causing migration problems of its own), comes from responsive vendors and all of your end users (including you) are happy with it, *please do not change it*. This is what I call the “All Pain, No Gain” migration. Even in the best cases, everyone will ask you, “Why did we do this?” In the worst cases, the migration will fail, you will be the goat, and your choice of Free Software will be held to blame.

Instead, look for new projects, or large projects using expensive hardware or software, or projects for which the software is *not* fulfilling their needs. This is where Free Software tends to be the most flexible and cost-effective solution. For new projects, the training costs are typically the same. There will be training costs for either proprietary software or Free Software, and this typically is not as much of a differentiator as it would be with re-training.

Thick and Thin

Down through the ages in computing we have moved from giant, single-program machines to

giant, multitasking machines to smaller, single-tasking mini-computers to multitasking mini-computers to smaller, single-tasking micro-computers to multitasking micro-computers and so forth, while still maintaining a lot of the older, “larger” computers. We also have moved from single mainframes to time-sharing systems to distributed systems and back again. In my view, what people really want is a time-sharing system of unlimited size and power, with very secure virtual firewalls, which can be available 25x8 (not just 24x7), and where backup and recovery are done automatically and with someone else’s money.

With the advent of the World Wide Web, a lot of applications now are going to be browser-based, with the applications and data (for the most part) residing on a back-office server. This promises an ease of system administration and security that are hard to supply with the pure distributed model.

Fortunately, the Linux Terminal Server Project (LTSP) solves a lot of the hard logistics of setting up a “thick and thin client” system. Although the costs of a modern-day desktop system reduce the need to squeeze every cent you can out of old hardware as desktop thin clients, it is still true that fat clients continue to expand in system resource needs while thin clients grow much more slowly, and are more stingy with desktop resources. It is also true that although hardware and networking have been increasing in capabilities over the years while prices have been dropping, the number of well-trained system administrators has not been keeping up, so a better model for end-user software configuration is needed.

Finding Applications

You can go several places to look for applications that meet your needs.

First, determine whether any of the applications you currently use and appreciate have gone “Free Software”. Many proprietary products now work on a free operating system or have developed a Free Software strategy, opening up their code and licensing while increasing their market share and support revenue. A very good example of this is Project.net (www.project.net), whose current owner, ICS, determined that making his project freely available and opening up the source code was the best way of doing business.

Other software may be developed directly from projects being listed on repositories, such as www.sourceforge.net and www.freshmeat.net. These repositories not only list the code and installation procedures, but also help weigh the receptiveness to the software from the community.

Finally, custom applications are not as expensive or difficult to build today as they were a few years ago. Using modern-day middleware, libraries of Free Software code and Web-based applications, you may find that developing an application tailored to your needs is a small investment compared to using an off-the-shelf application that requires you to change the way you do business to fit the application.

USER FRIENDLY by J.D. “Illid” Frazer



Cross-Over Applications

In some cases, stubborn applications keep people from moving to the environment they desire. Some of these applications are needed infrequently and may be handled by a dual-booting system or by running a product, such as VMware (www.vmware.com) or Win4Lin (www.win4lin.com), to allow you to run the applications simultaneously, albeit at a very slight performance hit.

Another great option is CodeWeavers' CrossOver products (www.codeweavers.com), now also available for the Intel OS X systems. CodeWeavers is based on the freely available Wine Project, and the parent company has been helpful in extending and expanding Wine's capabilities for many years.

Final Steps

Although I have heard of migrations that have gone cold turkey (turning off one system while turning the other on) successfully, I have heard of many more that failed. Nothing takes the place of a good transition strategy when going from an old system to a new one. Parallel running of the two systems is best, along with testing to see whether archived data is still available on the new systems.

Another good trick is to get the most enthusiastic office workers involved early and make sure that they have a good experience as they migrate over to the new tools. Every office has people like this. They buy the latest and greatest gadgets and are openly receptive to new things. Once they are enthusiastic about the new system, they often can help move other people over.

Last Recommendations

Do not be afraid to think outside of the box. I met a man with very old legacy code that was working well on very old hardware. He was concerned that the very old hardware was becoming more and more difficult to replace, and wanted to port it to Linux. I told him that although porting was a possibility, I would just run a hardware emulator for that hardware on top of Linux and use that to support his applications and customers "forever". He looked at me strangely, smiled, and walked away.

Likewise there are Free Software DOS emulators will allow DOS applications "forever", and modern-day CPU speeds sometimes make these applications run blazingly fast.

Use portable languages like Perl, Python and others to make your applications run on as many systems as possible.

Finally, when your system works well, evangelize what you have done. Write a paper about it, talk to your local Linux User Group, give a talk at LinuxWorld or write an article for *Linux Journal*. After all, probably a hundred or more other people are exactly where you are today and would like to have more freedom in their software.

—Jon "maddog" Hall

www.faircom.com/go/?track

Your packages get

tracked & delivered

on time by FedEx.

FairCom database technology makes it possible.



FairCom[®]

Tech Tips with Gnull and Voyd

Recover a dropped MySQL table and save partition images. CHESTER GNULL AND LAVERTA VOYD

Hey there sweeties, I'm Laverta Voyd, and my husband Chester picked some tips for y'all. Like I says last time, Chester ain't no talker, so I'm here to do the writin'. Now, I knows you ain't had much time to send in tips since last month, but we can't be doing this service for all you sweethearts out there without your help. So send in a tip, honey, so we can keep this steamboat rolling. Chester and me, we want this to keep afloat. We can use the money for hosting this here column, and you won't make out so bad yourself. This mag will send you \$100 for every tip you give us. So get the lead out of them undies of yours and send in some tech stuff Chester can use. He'll pick 'em and I'll sort it all out for him. In the meantime, we got that dear editor to stick in another tip, and we got a goodie from ol' Paddy. So here's y'are.

Recover a MySQL Table with Zmanda Recovery Manager

► *Somebody dropped a MySQL table. Duh. Fire your DBA. If you can't fire the DBA, then this tip helps.—Chester*

You are a MySQL database administrator. You take regular backups of your MySQL database. Somebody drops a table critical to the MySQL application (for example, the "accounts" table in a SugarCRM application). The MySQL application no longer works. How can you recover from the situation?

The answer is MySQL binary logs. Binary logs track all updates to the database with minimal impact on database performance. MySQL binary logs have to be enabled on the server. You can use the `mysqlbinlog` MySQL command to recover from the binary logs.

A better and more comprehensive solution is to use the Zmanda Recovery Manager (ZRM) for MySQL (MySQL backup and recovery manager). The `mysql-zrm` tool allows users to browse the binary logs and selectively restore the database from incremental backups:

```
# mysql-zrm --action parse-binlogs --source-directory=/var/lib
/mysql/sugarcrm/20060915101613
Log filename | Log Position | Timestamp | Event Type | Event
-----
/var/lib/mysql/my-bin.000015 | 11013 | 06-09-12 06:20:03 | Xid = 4413 | COMMIT;
/var/lib/mysql/my-bin.000015 | 11159 | 06-09-12 06:20:03 | Query | DROP TABLE IF EXISTS
`accounts`;
```

Here we're doing selective recovery for incremental backups without the DROP customer table from the SugarCRM database. Do two selective restore commands to restore from the incremental backup done on Sept 15, 2006, without executing the database event DROP TABLE at log position 11159:

```
# mysql-zrm --action restore --backup-set sugarcrm \
--source-directory=/var/lib/mysql/ sugarcrm/20060915101613/ \
--stop-position 11014
# mysql-zrm --action restore --backup-set sugarcrm \
--source-directory=/var/lib/mysql/ sugarcrm/20060915101613/ \
--start-position 11160
```

See the Zmanda Recovery Manager for MySQL for more information: mysqlbackup.zmanda.com.—Paddy Sreenivasan

Hedge Your Upgrade Bets by Using Partimage, Even on AMD64

► *Sometimes it don't make sense to assume upgrades are a good thing. Run partimage to save an image of your old version first. If the new distro version don't work, you can run partimage to put the old one back.—Chester*

I run multiple distros for various reasons. For example, I run 32-bit Kubuntu and 64-bit Kubuntu. I spend most of my time using 64-bit Kubuntu, but some packages are easier to set up and use on 32-bit Kubuntu, like Skype, for instance.

I recently decided to upgrade my 32-bit Kubuntu from Dapper to Edgy. Whenever you upgrade, you run the risk of breaking some existing programs. So I usually boot to another distribution and use partimage to save an image of the partition with the distribution I'm about to upgrade. If the upgrade doesn't go well, I always can use partimage to restore the partition to its previous state. There are many other reasons why you might want to save partition images, so you'll find this procedure useful even if you don't share my motivation.

Running Partimage on 64-bit Kubuntu AMD64

Partimage is finicky and refuses to run on a 64-bit system. For reasons beyond my knowledge, there is a partimage package you can install for Kubuntu AMD64, but it won't run. In my case, I want to boot to Kubuntu AMD64 and use partimage to save my 32-bit installation of Kubuntu, so the fact that partimage for Kubuntu AMD64 doesn't work is a major problem.

Or is it? It's actually quite easy to get partimage working on Kubuntu AMD64. Simply download the static binary form of partimage from www.partimage.org. Untar the binaries partimage and partimaged into /usr/sbin. These binaries should work fine even under AMD64. You should be able to type the command `partimage` at a root prompt to get it running. One would think it would be necessary to precede the command with `linux32`, but it works without it.


You should check to see if you have a /dev/dm inode before you use partimage. If you don't, you will be greeted with a screen like the



Figure 1. Partimage wants the /dev/dm inode created.

QUAD CORE. SAY NO MORE!

ABERDEEN STIRLING 128T



1U Quad-Core 4 SATA/SCSI/SAS 3 TB

- Up to two Quad-Core or Dual-Core Intel® Xeon® processors with up to 1333 FSB
- Intel® 5000P "Blackford" Chipset with 64-Bit Support
- Up to 64 GB 667/533 MHz Fully Buffered ECC DDR2 SDRAM
- Up to 4 x 750 GB (3.0 TB) Hot-Swap SATA Hard Drives
- Up to 4 SCSI/SAS Hard Drives available (Stirling 128S/128Q)
- 700 W AC Power Supply w/PFC
- **5-Year Warranty**

Starting at **\$1,999**

ABERDEEN STIRLING 228T



2U Quad-Core 8 SATA/SCSI/SAS 6 TB

- Up to two Quad-Core or Dual-Core Intel® Xeon® processors with up to 1333 FSB
- Intel® 5000P "Blackford" Chipset with 64-Bit Support
- Up to 64 GB 667/533 MHz Fully Buffered ECC DDR2 SDRAM
- Up to 8 x 750 GB (6.0 TB) Hot-Swap SATA Hard Drives
- Up to 8 SCSI/SAS Hard Drives available (Stirling 228S/228Q)
- 700 W High-Efficiency Redundant Power Supply
- **5-Year Warranty**

Starting at **\$2,429**

ABERDEEN STIRLING X314



3U Quad-Core 12 SATA 9 TB

- Up to two Quad-Core or Dual-Core Intel® Xeon® processors with up to 1333 FSB
- Intel® 5000P "Blackford" Chipset with 64-Bit Support
- Up to 32 GB 667/533 MHz Fully Buffered ECC DDR2 SDRAM
- Up to 12 x 750 GB (9.0 TB) Hot-Swap SATA Hard Drives
- 800+ MB/sec sustained data throughput RAID Controller
- 650 W Redundant Hot-Swap Power Supply
- **5-Year Warranty**

Starting at **\$3,849**

ABERDEEN STIRLING X418



4U Quad-Core 16 SATA 12 TB

- Up to two Quad-Core or Dual-Core Intel® Xeon® processors with up to 1333 FSB
- Intel® 5000P "Blackford" Chipset with 64-Bit Support
- Up to 32 GB 667/533 MHz Fully Buffered ECC DDR2 SDRAM
- Up to 16 x 750 GB (12.0 TB) Hot-Swap SATA Hard Drives
- 800+ MB/sec sustained data throughput RAID Controller
- 650 W Redundant Hot-Swap Power Supply
- **5-Year Warranty**

Starting at **\$4,449**

ABERDEEN STIRLING X526



5U Quad-Core 24 SATA 18 TB

- Up to two Quad-Core or Dual-Core Intel® Xeon® processors with up to 1333 FSB
- Intel® 5000P "Blackford" Chipset with 64-Bit Support
- Up to 32 GB 667/533 MHz Fully Buffered ECC DDR2 SDRAM
- Up to 24 x 750 GB (18.0 TB) Hot-Swap SATA Hard Drives
- 800+ MB/sec sustained data throughput RAID Controller
- 950 W Triple Redundant Hot-Swap Power Supply
- **5-Year Warranty**

Starting at **\$5,699**

ABERDEEN STIRLING X632



6U Quad-Core 32 SATA 24 TB

- Up to two Quad-Core or Dual-Core Intel® Xeon® processors with up to 1333 FSB
- Intel® 5000P "Blackford" Chipset with 64-Bit Support
- Up to 32 GB 667/533 MHz Fully Buffered ECC DDR2 SDRAM
- Up to 32 x 750 GB (34.0 TB) Hot-Swap SATA Hard Drives
- 800+ MB/sec sustained data throughput RAID Controller
- 1350 W Redundant Hot-Swap Power Supply
- **5-Year Warranty**

Starting at **\$7,399**

We couldn't put it any better than the words of Intel's Server Platform Group Marketing General Manager, Boyd Davis, when he said, "Aberdeen Stirling servers take advantage of the industry leading performance delivered by the new Quad-Core Intel Xeon processors to provide server and storage solutions that provide more performance, enable more efficiency and deliver more dependability." So if industry leading performance, more efficiency and dependability are on your server wish list, look no further. An Aberdeen Stirling server is the right choice for you. As Mr. Davis said, "Intel and Aberdeen both work to develop cutting-edge solutions for server customers."



"The Ultimate Linux Server... too fast for our benchmarks... we recommend the Aberdeen line of servers without reservation."

Linux Journal, August 2006—Aberdeen Stonehaven A261T

"AberNAS 128 boasts outstanding features"

Network Computing, April 2006—Aberdeen AberNAS 128

Finalist: Best Servers

LinuxWorld Product Excellence Awards, August 2005—Aberdeen Stonehaven A261

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

CRN Test Center Recommended, August 2005—Aberdeen Stonehaven A381

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

PC Magazine, February 2005—Aberdeen Stonehaven A261S

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

CRN Test Center Recommended, January 2005—Aberdeen AberSAN i100

[GNU AND VOYD]

one shown in Figure 1.

If nothing seems to happen when you press OK, partimage may not be able to create /dev/dm for you. Sometimes you can get past this screen by pressing OK several times, and partimage will work even if it doesn't create /dev/dm. Why take chances, though? Don't even start partimage until you create /dev/dm yourself with this command (assuming it's not already there):

```
# mknod -m 644 /dev/dm b 240 0
```

Now start partimage. Select the partition you want to save. In my case, this is partition /dev/sdb8. Then, enter a full path to the file where you want to store the image. I like to include the distro type, filesystem type and partition in the filename so I can remember why I created the partition image. The filename I used in this case is dapper32-ext3-sdb8.img. See Figure 2 for the example. The path /storage/disk/images points to a partition where I have lots of extra disk space for saving partition images.

Press function key F5 to continue. I use the default compression method, gzip, as it is much faster. If you are tight on disk space and don't mind waiting, you can choose bzip2 for the compression method. Partimage will check the partition and let you enter a description by default. I usually uncheck the description feature because the description is in the filename, but you may want to use the feature. Uncheck the feature by using the arrow keys to highlight it, and then press the spacebar to toggle to unchecked. You can see the screen with my choices in Figure 3.

Press F5 to continue. Partimage should get busy checking the filesystem.

Partimage can be flaky and sometimes reports errors that do not exist. If you have this problem, check the partition manually (with fsck.ext3, for

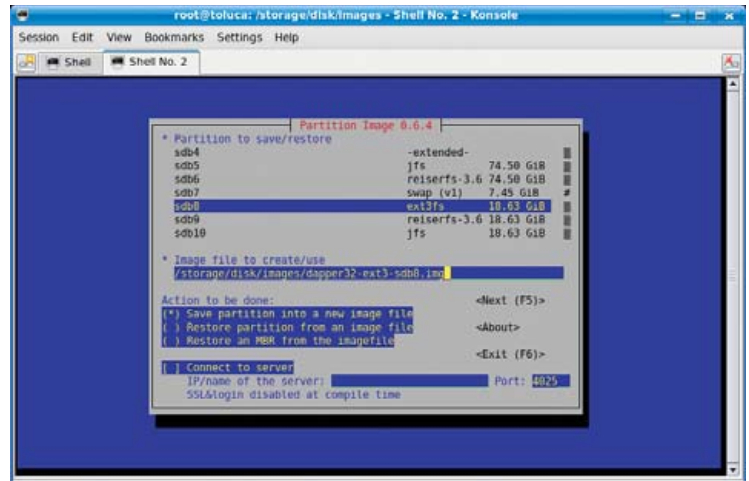


Figure 2. Save the image with a descriptive filename.

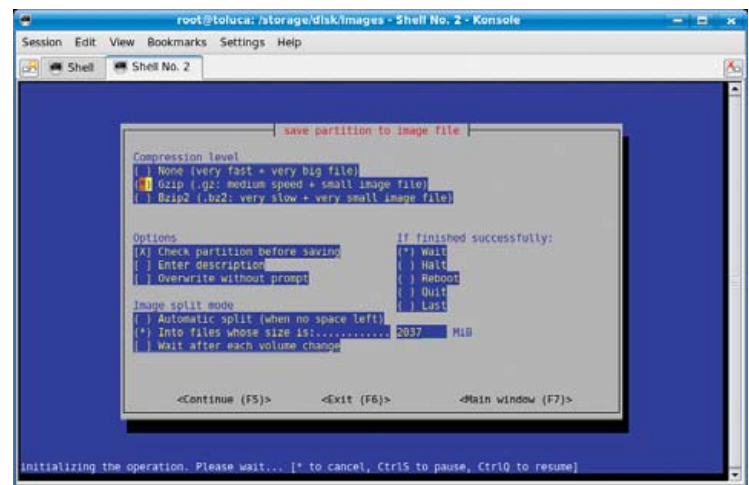


Figure 3. You can uncheck the description if you use a descriptive filename.

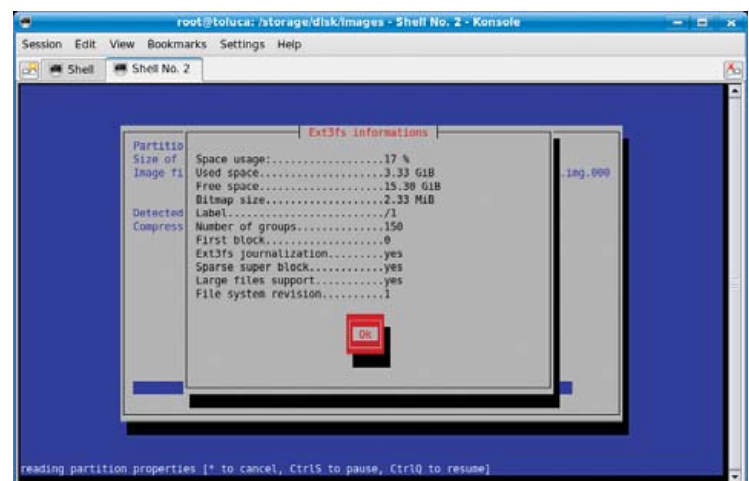



Figure 4. Partimage is ready to create an image for you.

FREE STUFF YOU ASK?

OK!

Send us a postage-paid, self-addressed envelope to the below address and we'll return a handful of Linux and Linux Journal stickers to you free of charge.



LINUX JOURNAL
Attn: Sticker promo
PO BOX 980985
Houston, TX 77098

example), and then uncheck the option to have partimage do it for you. If everything went well, you should see a screen something like the one shown in Figure 4.

Press Enter, and partimage will create one or more compressed image files of the partition. The rest of the process is self-explanatory.

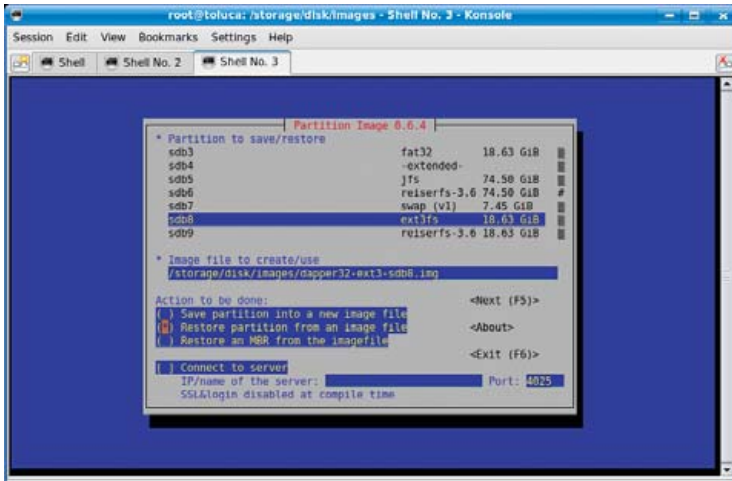


Figure 5. Restore the image to the partition.

Note that partimage adds numbers to your filename, because it prepares to split up the image into multiple files if necessary. Even if it needs only a single file, it will change the above filename to dapper32-ext3-sdb8.img.000.

It's now safe to boot into 32-bit Dapper and upgrade 32-bit Dapper to 32-bit Edgy. If everything works, you're in business.

If you run into problems and feel like you need to downgrade back to Dapper, all you have to do is boot back into 64-bit Kubuntu (or whatever other distro you were using) and run partimage again. This time, highlight the partition, type the same image filename (remember to add .000 to the name), but use the arrow keys to select Restore partition from an image file, and press the spacebar to select it. See Figure 5 for an example. Press F5 and the rest of the process should be self-explanatory. If partimage created multiple files ending in .000, .001 and so on, you don't have to worry about specifying them all. It will find the extra image files and restore them automatically.

Partimage works with many filesystems, but some of them have only beta support. Use partimage with filesystems like JFS with caution. It has worked for me, but that doesn't mean it will work for you. If you want to save and restore XFS filesystems, you should bypass partimage and use the XFS utilities xfsdump and xfsrestore, designed for saving and restoring partitions.—Nicholas Petreley

Linux Journal pays \$100 for any tech tips we publish. Send your tips with information on how you can be reached to techtips@linuxjournal.com. ■

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

**Flat Rate
Gigabit Ethernet**

1,000 Mbps of IP

\$13,000/month*

**Full 100 Mbps
Port**

Full Duplex

\$2,000/month


**Colocation Full
Cabinet**

Holds up to 42 1U
servers

\$400/month

Order Today!

email sales@he.net or call 510.580.4190



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



REUVEN M. LERNER

Scriptaculous

Scriptaculous is a spectacular set of libraries for Ajax.

Ajax is the hot new Web development paradigm that uses JavaScript to send and then handle asynchronous HTTP requests. The past few months, this column has looked into different ways to spawn and handle Ajax calls. The most complicated way was to do it ourselves, creating an XMLHttpRequest object, and then using it to send a request to the user's browser as well as to specify which JavaScript function will handle the response. Last month, we showed that we can simplify our lives greatly by using Prototype, a JavaScript library that includes many of the shortcuts and utility functions that are of use to JavaScript programmers.

The good news is that Prototype does indeed make JavaScript programming easier and more straightforward. But, one of the things people most want to do with JavaScript is create more flexible GUIs. This is especially true now that Web applications are becoming more desktop-like; users expect to have the same sense of feedback and control that they have with their nonbrowser applications.

Just as we saw with Ajax, there are ways to create and reuse these behaviors on your own. But why would you do that, when there are libraries that handle such tasks for you? Several of these are Scriptaculous, Moochikit and Dojo. (Dojo is actually a complete top-to-bottom JavaScript library; I expect to look at it in a future installment of this column.) This month, we look at Scriptaculous, an open-source library written by Thomas Fuchs. Scriptaculous makes it easy to spruce up our HTML files without having to delve into the low-level JavaScript.

Basics of Scriptaculous

Installing Scriptaculous couldn't be easier. Download the latest version of Scriptaculous (script.aculo.us), and install the six included JavaScript files (in the src directory) somewhere in your Web server's document root.

Actually, installing Scriptaculous could be even easier than this. If you use a recent version of Ruby on Rails, Scriptaculous and Prototype are already installed. See the Rails documentation for a description of how to use these libraries directly, as well as from Ruby code.

Note that Scriptaculous 1.6.5, which I use in this article, requires Prototype 1.5 or above. Although Prototype 1.5 likely will be released by the time this column sees print, I currently am relying on Prototype 1.5 RC1. Thus, there might be some differences between the functionality I describe here and the final distribution.

In order to use Scriptaculous, you need to include two script tags in your HTML page to load Prototype and then Scriptaculous:

```
<script src="/javascripts/prototype.js"
type="text/javascript"></script>
<script src="/javascripts/scriptaculous.js"
type="text/javascript"></script>
```

Simple Effects

So, what might you want to do with Scriptaculous? One of its most common uses is in the creation of visual effects. Each effect is defined as a method within the Effect object. You can create an effect by saying:

```
new Effect.EffectName('id')
```

where EffectName is the name of the effect that you want, and id is the ID of the HTML element on which the effect will take place. For example, if we have the following headline:

```
<h2 id="headline">The headline</h2>
```

we can make it fade by invoking:

```
new Effect.Fade('headline');
```

Of course, it makes sense for such things to happen only when particular events occur. Listing 1 contains a simple HTML file, effects.html, with two buttons labeled appear headline and fade headline. Clicking on the first button invokes Effect.Appear on the node with an ID of headline. Note that we don't pass the node itself to Effect.Fade, but rather the ID. Effect.Fade uses Prototype's \$() function to retrieve the node with that ID.

To make the headline fade, we set the following event

Listing 1. effects.html

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

  <head><title>Special effects</title>

  <script src="/javascripts/prototype.js"
type="text/javascript"></script>
  <script src="/javascripts/scriptaculous.js"
type="text/javascript"></script>
  </head>

  <body>
  <h2 id="headline">Welcome</h2>
  <p>Welcome to the page of effects!</p>
  <p><input type="button" value="Fade headline" onclick="new
Effect.Fade('headline')" /></p>
  <p><input type="button" value="Appear headline" onclick="new
Effect.Appear('headline')" /></p>

  </body>
</html>
```

handler:

```
onClick="new Effect.Fade('headline')"
```

Each of the effects has a number of settings, each of which is given a default value by Scriptaculous. To override one or more of these defaults, pass one or more of them in the invocation:

```
onClick="new Effect.Fade('headline', {delay: 2, duration: 10})"
```

In some cases, such as the appear/fade duo in Listing 1, it seems a bit silly to have two buttons. After all, what happens if I click on fade twice? It would be more reasonable to have a single button that turns the headline on when it's off and vice versa. Scriptaculous supports this with the toggle effect. For example, we can remove one button and give the second one an event handler that looks like this:

```
onClick="new Effect.toggle('headline', 'appear')"
```

Now, clicking on that button toggles the visibility of the headline, using appear and fade to turn the headline on

and off. We can do the same thing with the blind (BlindUp and BlindDown) and slide (SlideUp and SlideDown) effects as well. We also can combine a toggle with the parameters shown earlier:

```
onClick="new Effect.toggle('headline', 'appear', {delay: 2, duration: 10})"
```

Only a few effects come in pairs. Several are useful only for removing text, as in the following:

```
new Effect.Fold("headline")
```

Others are made to get the user's attention. For example, we can highlight our headline by doing this:

```
new Effect.Highlight("headline")
```

Autocomplete

Effects are certainly an important and impressive part of Scriptaculous; the library comes with many effects, and there are numerous ways to combine and invoke them. However, Scriptaculous offers much more than merely a bunch of



Expert Included.

Matt brings years of training and experience to developing innovative, powerful solutions.

He likes systems based on the Next-Generation AMD Opteron™ processor because AMD Virtualization™ allows safer testing of new operating systems and applications, with painless rollback until everything is just right. Matt also knows our customers appreciate the efficiency of AMD's Direct Connect Architecture.

When you partner with Silicon Mechanics, you get more than a next-generation AMD solution — you get an expert like Matt.



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

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

effects. It also includes some user-level functionality JavaScript programmers might want.

One such example, popularized by Google Suggest (and its cousin in Firefox 2.0) is an auto-completing text field. Such a text field lets you enter whatever contents

Scriptaculous makes it easy to spruce up our HTML files without having to delve into the low-level JavaScript.

you have entered so far matches a known text string, you are offered the chance to complete it. This is similar in many ways to the combo box widget that was long popular on Windows

systems, but which was unavailable to Web applications. (Some Scriptaculous users have created derivatives of the Autocompleter class that is more similar to a combo box.)

Scriptaculous comes with two different types of auto-completing text fields, differing only in how the completion list is filled. In the first case, known as `Autocompleter.Local`, the list of matches is set in JavaScript. The related text field, `Ajax.Autocompleter`, uses Ajax to retrieve a list of matches from a remote HTTP server. The two are similar enough in spirit that I demonstrate only `Autocompleter.Local` here for the sake of simplicity.

In order to use `Autocompleter.Local`, we first create a text field, just as we would do in an HTML form:

```
<input type="text" id="distro" name="distro_text_field"
autocomplete="off" />
```

Listing 2. complete.html

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

  <head><title>Special effects</title>

  <script src="/javascripts/prototype.js"
type="text/javascript"></script>
  <script src="/javascripts/scriptaculous.js"
type="text/javascript"></script>
  </head>

  <body>
    <h2>Enter your favorite distribution</h2>
    <p><input type="text" id="distro" name="distro_text_field"
autocomplete="off" /></p>
    <div class="auto_complete" id="distro_complete"
style="display:none"></div>

    <script type="text/javascript">
      new Autocompleter.Local('distro', 'distro_complete', ['Red Hat',
'Fedora Core', 'Debian', 'Gentoo', 'Knoppix', 'Ubuntu', 'Kubuntu'], { });
    </script>

  </body>
</html>
```

Notice that I have included the setting `autocomplete="off"` in the above field. The `autocomplete` attribute, which is used only by Microsoft's Internet Explorer, is set here to off to ensure that IE users aren't unpleasantly surprised by dueling completion systems.

Next, we create a `<div>` section with an ID attribute. In addition, we set the `style` attribute to keep the div hidden until we modify its styling:

```
<div id="distro_complete" style="display:none"></div>
```

Finally, we add some JavaScript that creates a new `Autocompleter.Local` object:

```
<script type="text/javascript">
  new Autocompleter.Local('distro', 'distro_complete',
  ['Red Hat', 'Fedora Core', 'Debian', 'Gentoo',
  'Knoppix', 'Ubuntu', 'Kubuntu'], { });
</script>
```

The constructor for `Autocompleter.Local` takes four arguments: the ID of the text field, the ID of the div into which we'll insert completions, an array containing the completion strings and a set of options (currently empty). If you try to put this code in the `<head>` of your document, it will fail with odd errors, because the text field and div must exist before the code is executed.

By including the above in an HTML page (as in Listing 2), you set the stage for auto-completion to work. Whenever the user loads the page and types a letter into the text field, Scriptaculous waits for 0.4 seconds of inactivity. If the user isn't typing, and if the text field already contains one or more characters, `Autocompleter.Local` tries to find a match from the current list. If it finds one, it fills in the rest of the text field.

If it finds more than one (as would happen to a user typing K in our example, which matches both Kubuntu and Knoppix), the system displays a menu of options, from which the user may choose by typing or clicking.

Conclusion

Prototype is a library aimed at the JavaScript that endeavors to solve many programmers' needs. However, Prototype's functionality extends only so far as the programmer; it doesn't offer any direct GUI improvements. Given that JavaScript is largely used to handle the GUI in a Web application, it shouldn't come as a surprise that there are several libraries built on top of Prototype to handle such tasks. Scriptaculous appears to be one of the best known of these, and in my experience, there's good reason for that.

There are several functions of Scriptaculous that we didn't get to explore this month, including drag and drop and JavaScript unit testing. This last item is probably of note even for JavaScript programmers who have no intention of creating GUI effects. ■

Resources for this article: www.linuxjournal.com/article/9505.

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

FOCUS ON WHAT MATTERS MOST

GET DUAL-CORE INTEL XEON PROCESSORS IN SERVERSDIRECT SYSTEMS AND CONCENTRATE ON BUSINESS OPPORTUNITIES INSTEAD OF SERVER HEADACHES



STARTING PRICE
\$2,499

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

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

FEATURE HIGHLIGHTS:

Intel Xeon 5050 3.0GHz Dual Core Processor (Dual Processor Option)

3U Chassis with 800W Redundant Power Supply

Intel® 5000P (Blackford) Chipset

Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2pcs x 512MB)

Seagate 400GB SATA-II 16MB Cache 7200RPM

16 x 1" Hot-swap SATA Drive Bays

ATI ES1000 Graphics with 16MB video memory

Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller

RAID 0, 1, 5, 10 Support

1U ENTRY LEVEL SERVER

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



\$999

1U Rackmount Chassis with 300W Power Supply

Intel® Pentium® D 930 3.0GHz Dual Core 800FSB

Intel® E7230 (Mukilteo) Chipset

Kingston 1024MB 667MHz DDR2 ECC (2pcs x 512MB)

Seagate 250GB SATA-II 7200RPM hard drive

2 x 1" Hot-swap SATA Drive Bays

ATI RageXL graphics, 8MB

2x Intel® 82573V (Tekoa) PCI-e Gigabit LAN Ports

RAID 0, 1

SKU# SDR-1110T

COST EFFECTIVE 2U SERVER

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



\$1,599

Intel Xeon 5050 3.0GHz Dual Core Processor (Dual Processor Option)

2U Chassis with 600W Power Supply

Intel® 5000V (Blackford VS) Chipset

Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2pcs x 512MB)

Seagate 400GB SATA-II 16MB Cache 7200RPM

6 x 1" Hot-swap SATA-II Drive Bays

ATI ES1000 Graphics with 16MB video memory

Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller

RAID 0, 1, 5, 10 support

SKU# SDR-2501T

4U TOWER/RACKMOUNT 64-BIT XEON ENTERPRISE SERVER

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



\$3,999

Intel Xeon 5050 3.0GHz Dual Core Processor (Dual Processor Option)

4U Rackmountable / Tower

Intel® 5000X (Greencreek) Chipset

Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2pcs x 512MB)

8pcs x Seagate 400GB SATA-II 16MB Cache 7200RPM

3Ware 9550SX-8port RAID Controller Card

GeForce 6200 TC 128MB OB 256M support DDRTV-out PCI-E

Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller

SKU# SDP-7045A-TB

SDR-5500T 5U ADVANCED STORAGE SERVER

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



\$6,799

Intel Xeon 5050 3.0GHz Dual Core Processor (Dual Processor Option)

5U Chassis with 24 hot-swap bays & 950W redundant power supply

Intel® 5000P (Blackford) Chipset

Kingston 1024MB 667MHz DDR2 ECC FB-DIMM (2pcs x 512MB)

12pcs x Western Digital 400GB SATA RAID Drive

1pc x 3Ware 9550SX-12 port RAID Controller Card

ATI ES1000 Graphics with 16MB video memory

Intel® (ESB2/Gilgal) 82563EB Dual-port Gigabit Ethernet Controller

RAID 0, 1, 5, 10 support

SKU# SDR-5500T

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





MARCEL GAGNÉ

Words, Words, Words...

When it comes to being understood and sharing information, it's not just about open source, it's about open standards.

What are you doing, François? Our guests will be here any moment, and you are still sitting in front of your computer. *Quoi?* Yes, I agree with you that it's a good idea to store all these old documents in OpenDocument format. I admire your desire to ensure the long-term usability of these documents by converting them from their limited, proprietary format, but this is not the way. We have thousands of documents from hundreds of people on this storage area network. Converting them one at a time as you are doing will take forever, and we are minutes away from opening time. Besides, I have a much better way to deal with this and you'll see it on tonight's menu.

Vite! To the wine cellar. I see our guests coming to the door now. There are six cases of 2002 Paso Robles Zinfandel over in the East wing, right next to the old door marked *Danger*—I should really have you check that out sometime so we can find out what is back there—bring the wine and I will greet our guests. *Vite!*

Welcome, *mes amis*, to *Chez Marcel*, where fine Linux and open-source fare is married with some of the world's best wines. Your tables are ready and waiting, so please, sit down and make yourselves comfortable. My faithful waiter, François, will return shortly from the cellar with your wine for tonight. Before you arrived, we were discussing a little project to convert all of the old proprietary format .doc documents to the OpenDocument format, OpenOffice.org's default document format. This is the OASIS OpenDocument XML (eXtensible Markup Language) format, an open standard for document formats (it is saved with an .odt extension). The OpenDocument format is the closest thing to document freedom you will get (short of plain text). The format is vendor and application-neutral. You are guaranteed support and portability because it is an open standard. Many organizations, such as the European Commission and the state of Massachusetts, are starting to recommend the OASIS OpenDocument format for the very reasons I've mentioned.

Ah, François, good to see you made it back with the wine. Please, pour for our guests. Enjoy, *mes amis*. You'll find this particular wine rich and jammy, with wonderful black raspberry flavors, a little licorice, a little pepper...

Ah, where was I? Oh, yes—converting to OpenDocument makes sense, but some people, of course, will stay with the Word format, not so much for technical reasons, as for simple inertia. After all, Microsoft Word is everywhere. The sheer number of Word installations is the very reason that OpenOffice.org was designed to support the Microsoft Office format as thoroughly as it does. That said, if you do want to switch to the OASIS OpenDocument format, OpenOffice.org Writer provides an easy way to do that. Rather than converting documents one by one, the Document Converter speeds up the process by allowing you to run all the documents in a specific directory in one pass. It also works in both directions, meaning you can

convert from Word to OpenOffice.org format and vice versa. The conversion creates a new file but leaves the original as it is. Here's how you do it.

From the menu bar, select File, move your mouse to Wizards, then select Document Converter from the submenu. To convert your Microsoft Office documents, click the Microsoft Office radio button, then check off the types of documents you want (Figure 1). You can do Excel and PowerPoint documents at the same time.



Figure 1. Start by selecting the types of documents you want to convert.

The next screen asks whether you want both documents and templates or just one or the other. You then type in the name of the directory you want to import from and save to. This can be the same directory or you can choose an alternate. You need to answer this set of questions three times if you've chosen to do Excel and PowerPoint files at the same time, but



Figure 2. If you choose to convert Excel and PowerPoint documents as well, you'll get a similar dialog for each one.

THE PENGUIN

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

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

MAY CAUSE DROWSINESS



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

www.penguincomputing.com



REALLIFELINUX

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

the dialog is the same as the one you'll see for Word documents (Figure 2).

After you've entered your information and gone to the next screen, the program confirms your choices and gives you a final chance to change your mind. Click Convert to continue. As the converter does its job, it lists the various files that it encounters and keeps track of the process. Click the Show Log button to see a listing of everything the converter encountered (Figure 3). When the job is done, you'll have a number of files with .odt extensions in your directory. Spreadsheets will have .ods extensions, and presentations will have .odp extensions. If you change your mind, don't worry. Your original files are still there, so you've lost nothing.



Figure 3. The progress dialog shows the status of the conversion and provides a log of the process.

As you can see, it's easy. And, this wine is easy-drinking, I see. François, some of our guests' glasses are looking a little empty. Please, *resolve* this issue for them with a little top-up. *Merci, mon ami.*

If you've never taken a good look at an OpenDocument document, you should. It's quite fascinating, actually. What you may not know is that the .odt file is actually a compressed file containing all the elements that make up your document. To be exact, it's a ZIP file. Let's say you had a document titled mydocument.odt that contained several images in addition to the text itself. To extract and view the elements, type the following in a shell or terminal window (you may want to do this in a temporary folder somewhere):

```
zip mydocument.odt
```

The result would look like this.

```
Archive:  mydocument.odt
 Length   Date    Time    Name
-----
      39  10-13-06  20:09  mimetype
      0  10-13-06  20:09  Configurations2/statusbar/
      0  10-13-06  20:09  Configurations2/accelerator/current.xml
      0  10-13-06  20:09  Configurations2/floater/
      0  10-13-06  20:09  Configurations2/popupmenu/
      0  10-13-06  20:09  Configurations2/progressbar/
      0  10-13-06  20:09  Configurations2/menubar/
      0  10-13-06  20:09  Configurations2/toolbar/
      0  10-13-06  20:09  Configurations2/images/Bitmaps/
 24634  10-13-06  20:09  Pictures/1235696243C.png
```

```
14808  10-13-06  20:09  Pictures/10C3F082746.png
68331  10-13-06  20:09  Pictures/20963618D3B.png
 1925  10-13-06  20:09  Pictures/19C4B78A82D.png
 9677  10-13-06  20:09  Pictures/112FEC43498.png
 6100  10-13-06  20:09  Pictures/1005A594DCB.png
172170  10-13-06  20:09  Pictures/3009CCB23C4.png
   54  10-13-06  20:09  layout-cache
23674  10-13-06  20:09  content.xml
 7950  10-13-06  20:09  styles.xml
 1211  10-13-06  20:09  meta.xml
 4899  10-13-06  20:09  Thumbnails/thumbnail.png
 7386  10-13-06  20:09  settings.xml
 2904  10-13-06  20:09  META-INF/manifest.xml
```

```
-----
345762                23 files
```

This collection of XML definitions, images and so on, makes the document portable and readable by other programs.

Of course, *mes amis*, even if you have all these old documents and you want to preserve them in some kind of open format that doesn't require a copy of Microsoft Office, you may not need them in editable format. A simple, read-only format, such as PDF, may be the answer. OpenOffice.org has a built-in export to PDF, but unlike the document converter, this is a one-at-a-time affair. As François can tell you, one at a time can take a long time.

If you have OpenOffice.org on your system, I have just the thing for you. It's an OpenOffice.org macro document called—wait for it—Document Converter, or just DocConverter. This macro, written by Danny Brewer and Don Horwood, is designed to let you do batch conversions of any document format OpenOffice.org supports to any other format it supports easily. In other words, the output doesn't have to be PDF, as there are a number of alternatives you can choose. You can find the Document Converter at the OpenOffice.org Macros Web site (see the on-line Resources). Macros are sorted into end-user applications and those suited to developers. Click the For End-Users link at the top of the page, and scroll down to find

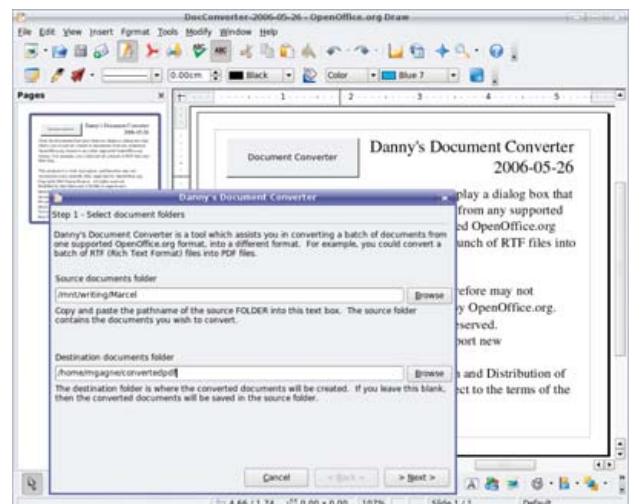


Figure 4. To make the document converter do its thing, simply click the big button and follow the instructions in the wizard.

Document Converter.

To use the macro, unzip the file and save the document somewhere. When you open it with OpenOffice.org Writer, a warning dialog appears, asking you if you want to enable the macros in the document. The correct answer, in this case, is yes. The document that appears is exactly that, a document. At the top left of the document is a big button labeled Document Converter (Figure 4). Click that button, then simply follow the wizard that pops up. Tell it which folder your Word files are in and what folder you would like the PDF files to appear in. It's a simple point-and-click task.

Don't forget to check out some of the other great macros that exist on the site. You might discover something you can't live without.

All this converting of documents using OpenOffice.org is pretty cool, but it also tends to make us forget the great and powerful conversion tools that lie just beneath the graphical surface of your Linux system. Most distributions come with a variety of document converters waiting for the command-line user to use them. For instance, you may have PostScript documents that you want to convert to PDF, so you can send them to friends or family who don't understand PostScript. The command-line program, `ps2pdf`, comes in extremely handy under those circumstances:

```
ps2pdf mydocument.ps mydocument.pdf
```

The `ps2pdf` program produces a document compatible with Acrobat Reader, version 3, also known as PDF, version 1.2. To create version 1.3 PDF output (for Acrobat Reader 4 or later), use `ps2pdf13`. There's also a `ps2pdf14` program. I'll leave it to you to guess which version of PDF it outputs. You also can convert PDF documents to PostScript by using `pdf2ps` and PostScript documents to plain ASCII text with `ps2text`. You'll also find a program called `ps2ascii`, which does more or less the same thing, but it doesn't handle encoded text (such as French accents) as well.

Hey, how about a nice, plain-text document from that Web site, minus all the HTML tags? That's the idea behind the `html2text` program. To define the output file, you need to specify it using the `-o` option:

```
html2text -o outputfile.txt http://somedomain.dom/document.html
```

If you are curious to see what sorts of conversions you can do, change directory to `/usr/bin` and look for the commands that include a 2 or a to. Not everything you see will be a document converter, of course, but you'll discover some interesting commands that are.

Before I leave the subject of Word document conversion completely, I need to mention Dom Lachowicz's `wvWare` (which started life as just `wv` when Caolón McNamara wrote it). The package is available from SourceForge (see Resources), but you should have no trouble finding a package for your particular Linux distribution. For `wv`, think "Word Viewer". This package allows you to convert (or view) Microsoft Word documents to a wide variety of formats. `wvWare` is actually a collection of command-line tools, such as `wvText`:

```
wvText SomeWordDocument.doc
```

The output will go directly to your screen, so you may want to capture it by redirecting to a file or piping it to more (or less). There's also `wvPDF` to convert to PDF, `wvLatex` to convert to LaTeX, `wvAbw` to create Abiword-compatible documents and more. Check out the site documentation for all the alternatives.

Why use all these text tools when the graphical alternative exists? The answer, *mes amis*, is speed. Speed *and* flexibility. Sorry, the two answers are speed and flexibility—okay, I'll stop right there.

All through this discussion, I've concentrated on text, but convert-

ing to open standards covers a lot of possibilities, including graphics, video files, music files and more. Tackling these formats is the beginning of another, rather rich, menu, but alas, that insistent clock is telling us that closing time is here. As you can see, there are many opportunities for taking those old, closed-format documents and storing them in formats you will be able to access years from now, free from the will or whim of some mega-corporation's definition (or support) of what it calls a standard. Plain text, *mes amis*, is still the most portable of all formats. Nevertheless, choosing and using an open document format, such as OpenDocument, allows you to take advantage of the portability of plain text *and* the richness of graphics and other non-text elements.

François, please refill our guests' glasses one more time. And now, *mes amis*, raise your glasses and let us all drink to one another's health. *A votre santé! Bon appétit!* ■

Resources for this article: www.linuxjournal.com/article/9509.

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

UNIX and Linux Performance Tuning Simplified!

Understand Exactly What's Happening	Maintain Full Control	Plan for Future Growth
SarCheck® translates pages of sar and ps output into a plain English or HTML report, complete with recommendations.	SarCheck fully explains each of its recommendations, providing the information needed to take intelligent informed actions.	SarCheck's Capacity Planning feature helps you to plan for growth, before slow downs or problems occur.

Make Your System Fly With SarCheck!

APTITUDE CORPORATION

Request your free demo at www.sarcheck.com

IBM Business Partner | Sun Partner Advantage | HP Business Partner | Linux



DAVE TAYLOR

Analyzing Your Search Keywords

Screen the unwanted results out of your access log searches.

Last month, we started exploring how you can use a shell script to extract and analyze the HTTP_REFERER values out of your Web server log and identify the most common terms and phrases that people used to find your pages. Sounds useful, doesn't it?

The problem is, the script is more nuanced than it initially seems. Last month, we wrapped up with the following shell script:

```
#!/bin/sh

ACCESSLOG="/var/logs/httpd.logs/access_log"

grep 'google.com/search' $ACCESSLOG | \
awk '{print $11}' | \
cut -d\? -f2 | cut -d\& -f1 | \
sed 's/+/ /g;s/%22/" /g;s/q=/' | \
sort | \
uniq -c | \
sort -rn | \
head -5
```

When I run this, here's what I see:

```
$ sh google-searches.sh
94 hl=en
18 client=safari
6 client=firefox-a
4 sourceid=navclient
4 client=opera
```

That's weird, because it's not search terms, it's other variables that are included with search strings sent from sites like Google (hl=en says that you've constrained searches to English-language sites only, client=safari identifies the user's Web browser as Apple's Safari and so on).

Screening Out False Matches

The problem is revealed when we look at the first ten matches rather than only the first five:

```
$ sh google-searches.sh | head -10
94 hl=en
18 client=safari
6 client=firefox-a
4 sourceid=navclient
4 client=opera
3 wicked cool scripts
3 hl=zh-CN
2 num=100
2 hs=wNy
2 barbara nelson%2Bpurses
```

Ah, so we can see that there are two valid searches here, one for "wicked cool scripts" and one for "Barbara nelson%2Bpurses". Not sure what the latter one's about, but it's useful and important to see. Fortunately, screening out the bogus matches is as simple as using grep to remove fields that include an equal sign: grep -v "=".

Rather than have that at the very end of the long pipe in the script, however, I'll place it immediately after the sed invocation to strip out false results as soon as possible in the pipeline to speed up the entire script. Now it looks like this:

```
grep 'google.com/search' $ACCESSLOG | \
awk '{print $11}' | \
cut -d\? -f2 | cut -d\& -f1 | \
sed 's/+/ /g;s/%22/" /g;s/q=/' | \
grep -v '=' | \
sort | \
uniq -c | \
sort -rn
```

Notice that the sed statement itself strips out the name= part of the search (q=), so that it's not incorrectly matched in the new grep statement.

Now we have the results we want:

```
$ sh google-searches.sh | head -10
3 wicked cool scripts
2 barbara nelson%2Bpurses
1 wsj%20password
1 why did animal kingdom introduce expedition everest
1 what makes a great speaker%3F
1 university of phoenix center of writing excellence
1 ubuntu x problem
1 triboot osx ubuntu ydl
1 the best dvd players
1 symbol html heart
```

This site doesn't get a huge amount of traffic, so let's run the very same script against my far-busier AskDaveTaylor.com site. The results are more interesting:

```
$ sh google-searches.sh | head -10
5 standalone player
4 psp help
4 create a myspace
4 Documents and Settings"
4 %24NtUninstall
3 view myspace accounts that are set to private
3 i cant hear music on runescape
2 transfer files to psp
2 sync v3 motorola mac
```



PerformanceWare 1550

- Dual Intel Dual or Quad Core Xeon Processors with up to 8MB Cache
- Intel EM64T Technology (64bit Extensions)
- Intel 5000P Chipset with 1333Mhz FSB
- Up to 16GB of PC2700 ECC Registered DDR Memory
- LSI 8 Port SAS RAID Controller
- Up to 584GB High Performance SAS Redundant Storage Capacity (438GB Raw)
- Onboard Dual Intel 10/100/1000 Ethernet Adapters
- Slim CD-ROM
- Slim Floppy Drive
- CentOS Linux
- 3 Year Return-to-Depot Warranty Service
- Optional Onsite Service Available

Starting at \$2899

StorageWare SA108

- Dual Intel Dual or Quad Core Xeon Processors with up to 8MB Cache
- Intel EM64T Technology (64bit Extensions)
- Intel 5000P Chipset with 1333Mhz FSB
- Up to 16GB of PC2700 ECC Registered DDR Memory
- 3Ware 4-Port Serial ATA RAID Card (Level 5 or optional RAID 6)
- Up to 2.25TB Redundant Storage Capacity (3TB Raw)
- Onboard Dual Intel 10/100/1000 Ethernet Adapters
- Slim CD-ROM
- Slim Floppy Drive
- CentOS Linux
- 3 Year Return-to-Depot Warranty Service
- Optional Onsite Service Available

Starting at \$2499

StorageWare SA360i

- Front Panel LCD Screen
- Supports the latest Intel Dual and Quad Core Xeon Processors
- Intel 5000P Chipset with 1333Mhz FSB
- Up to 16GB of 667MHz ECC Registered DDR2 Memory
- AMCC/3Ware 16-Port Serial ATA RAID Card with optional Battery Backup Unit
- LSI External SAS RAID Controller
- Up to 11.25TB of Native Redundant Storage Capacity (12TB Raw)
- Onboard Dual 10/100/1000 Ethernet
- Slim CD-RW/DVD-ROM
- CentOS 4.2 Linux
- 3 Year Return-to-Depot Warranty Service

Starting at \$8999

Maximum data protection with unaffected performance.

Introducing RAID 6 into our family of robust systems that fits your every need.



StorageWare DA8

- Front Panel LCD Screen
- Up to 5.25TB of available redundant RAID 5 Storage
- 256MB of onboard Cache per controller
- Two 2GBps Fibre or SCSI links

- 3 Year Return-to-Depot Warranty Service
- Optional Onsite Service Available

Starting at \$5649



StorageWare DA12

- Up to 8.25TB of available redundant RAID 5 Storage
- 256MB of onboard Cache per controller
- Two 2GBps Fibre or SCSI links
- Optional 4GBps SATA to FC link

- 3 Year Return-to-Depot Warranty Service
- Optional Onsite Service Available

Starting at \$8499

Each DA system can connect to as many as seven external JBOD units for a total of up to 96TB Available RAID 5 Storage.



StorageWare DA16

- Up to 11.25TB of available redundant RAID 5 Storage
- Up to 512MB of onboard Cache per controller
- 2GBps Fibre or SCSI links

- Optional 4GBps SATA to FC link
- 3 Year Return-to-Depot Warranty Service
- Optional Onsite Service Available

Starting at \$8999



StorageWare DA24

- Up to 17.25TB of available redundant RAID 5 Storage
- 512MB of onboard Cache per controller
- 2GBps Fibre or SCSI links

- 3 Year Return-to-Depot Warranty Service
- Optional Onsite Service Available


Starting at \$11,299



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

To get started, contact us at 888.828.POGO or inquiries04@pogolinux.com
Pogo Linux, Inc. 701 Fifth Ave. Suite 6850, Seattle, WA 98104

www.pogolinux.com



```
2 running unix in windows xp
```

Much more interesting. Oh, and if you want to know how many searches you're exploring, it's another simple tweak to the script, an invocation of `wc`:

```
$ sh google-searches.sh | wc -l
501
```

So out of 501 searches, the single-most common search is "standalone player", which represents only five out of 500, or 1% of my search traffic.

Stripping Out Unwanted Characters

One more step before we walk away from this script for the month: let's get rid of the strange characters that have been carried over from the original URL encoding of the user's Web browser. What am I talking about? The `%24`, the closing double quote in Documents and Settings and the `%2B` in the earlier search for purses.

You can figure out all the mappings and convert everything as appropriate, but I'm lazy at the end of the day and will instead simply find all `%xx` sequences and replace them with a single space.

This sounds hard, but it's a perfect job for `sed` because it

allows you to do pattern matching and then replace the matched material with whatever else you desire. Here's how I'd do that:

```
sed 's/%[0-9a-fA-F][0-9a-fA-F]/ /g'
```

Let's look at this closely before you panic. A set separated by square brackets is a set in regular expression terminology, so `[0-9]` will match any of 0, 1, 2, 3, 4, 5, 6, 7, 8, 9 or 0. It turns out that URL encoding uses hexadecimal, so not only can the values be 0-9, but they can also be A, B, C, D, E and F, in upper- or lowercase letters—hence 0-9 and a-f and A-F in the pattern. The overall pattern form is `%` followed by any of these possible values, followed by any of these possible values again. Now you can see the full pattern.

Finally, before we beat this completely into the ground, note that the bigger structure here in the `sed` statement is `s/old/new/g`, which replaces old with new throughout the line, whether it occurs once or 15 times.

We're not quite done yet, however, because we also need to strip the stray double quotes. Again, this is easily added to the `sed` statement:

```
sed 's/%[0-9a-fA-F][0-9a-fA-F]/ /g;s/"//g'
```

Here's the final script:

```
grep 'google.com/search' $ACCESSLOG | \
awk '{print $11}' | \
cut -d\? -f2 | cut -d\& -f1 | \
sed 's/+//g;s/%22/"g;s/q=/' | \
sed 's/%[0-9a-fA-F][0-9a-fA-F]/ /g;s/"//g' | \
grep -v '=' | sort | uniq -c | sort -rn
```

And, the final results:

```
$ sh google-searches.sh | head -15
6 NtUninstall
5 standalone player
4 psp music
4 psp help
4 creat a myspace
4 Documents and Settings
3 view myspace accounts that are set to private
3 i cant hear music on runescape
2 transfer files to psp
2 sync v3 motorola mac
2 running unix in windows xp
2 rss feed reader shell
2 reinstall windows xp hp
2 psp transfer music
2 psp internet
```

Note that once we remove the stray material, things organize slightly differently (for example, here you can see that `psp music` is one of the top searches, but earlier we had different variations of `psp music`, and it didn't make it to a top search value).

Okay, enough torturing of the Apache log file. Let's wrap this up and we'll switch to something completely different in the next column! Suggestions? Please e-mail them to me! ■

Dave Taylor is a 26-year veteran of UNIX, creator of The Elm Mail System, and most recently author of both the best-selling *Wicked Cool Shell Scripts* and *Teach Yourself Unix in 24 Hours*, among his 16 technical books. His main Web site is at www.intuitive.com.

Ultra Dense, Powerful, Reliable... Datacenter Management Simplified!

15" Deep, 2-Xeon/Opteron or P4 (w/RAID) options



Customized Solutions for... Linux, BSD, W2K

High Performance Networking Solutions

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

Rackmount Server Products

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



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

CALL: 1-800-921-IRON

SUPERMICRO®

Xeon 5300/5100/5000 Server Solutions

• Multiple-Core • 64GB FBD Memory

• High Efficiency Power



1U SuperServer 6015B-8+

- ✓ Dual Intel® Quad-core/Dual-core Xeon® 5300/5100/5000 sequence
- ✓ SUPER® X7DBR-8+ serverboard
- ✓ Intel 5000P Chipset/1333MHz FSB
- ✓ Up to 64GB fully-buffered DIMM (FBD)
- ✓ 1 PCI-X 133MHz or 1 PCI-E x16 or 1 PCI-E x8
- ✓ Adaptec dual-channel U320 SCSI
- ✓ AOC-LPZCR2 (Zero-channel RAID) support
- ✓ Dual Gigabit LAN & 16MB graphics
- ✓ 4 hot-swap SCA drive bays w/SAF-TE
- ✓ 1 slim floppy & 1 slim DVD-ROM
- ✓ 700W high-efficiency power supply w/l²C
- ✓ IPMI 2.0 with KVM-Over-LAN support
- ✓ IPMI 2.0 with virtual media over LAN & optional KVM-over-LAN support

2U SuperServer 6025B-8R+

- ✓ Dual Intel® Quad-core/Dual-core Xeon® 5300/5100/5000 sequence
- ✓ SUPER® X7DB8+ serverboard
- ✓ Intel 5000P Chipset/1333MHz FSB
- ✓ Up to 64GB fully-buffered DIMM (FBD)
- ✓ 3 PCI-X 133/100MHz, 2 PCI-E x8 & 1 PCI-E x4
- ✓ Adaptec dual-channel U320 SCSI
- ✓ AOC-LPZCR2 (Zero-channel RAID) support
- ✓ Dual Gigabit LAN & 16MB graphics
- ✓ 8 hot-swap SCA drive bays w/SAF-TE
- ✓ 1 slim DVD-ROM, optional USB ports & FDD
- ✓ 700W redundant high-efficiency power supply w/l²C
- ✓ IPMI 2.0 with virtual media over LAN & optional KVM-over-LAN support

4U SuperServer 7045B-3

- ✓ Dual Intel® Quad-core/Dual-core Xeon® 5300/5100/5000 sequence
- ✓ SUPER® X7DB3 serverboard
- ✓ Intel 5000P Chipset/1333MHz
- ✓ Up to 32GB fully-buffered DIMM (FBD)
- ✓ 3 PCI-X 133/100MHz, 2 PCI-E x8 & 1 PCI-E x4
- ✓ Adaptec controller for 8 SAS/SATA ports
- ✓ AOC-LPZCR2 (Zero-channel RAID) support
- ✓ Dual Gigabit LAN & 16MB graphics
- ✓ 8 hot-swap SAS drive bays w/SES2
- ✓ 90° rotatable module: USB ports, FDD, 2 drive bays
- ✓ 100% cooling redundancy: 6 fans & air shroud
- ✓ 650W power supply with redundant cooling
- ✓ IPMI 2.0 with virtual media over LAN & optional KVM-over-LAN support

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

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

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

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

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

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

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

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



MICK BAUER

Introduction to SELinux

Invest some time into SELinux and worry less about zero-day attacks.

SELinux, the NSA's powerful implementation of mandatory access controls for Linux, can seem like a daunting technology. It's got a lot of moving parts that are labeled (pun intended) with arcane, acronym-intensive terminology, adding some very dense layers of abstraction over Linux's already-abstract architecture. To compound the problem, much of SELinux's documentation seems to have been written by security geeks for security geeks.

Well, people say all that and worse about LDAP too, but as with LDAP (which we covered in this column in the July, August and September 2003 issues of *LJ*), you can make SELinux do what you need it to do if you learn some basic concepts, become familiar with a modestly sized list of terms and study some representative policy files.

In this month's column, we discuss SELinux basics. We begin with SELinux's general design goal; introduce the concepts of SELinux subjects, permissions and objects, and how they fit into security contexts; and tie those ideas together in a discussion of Type Enforcement.

Believe me, that's plenty to start off with! We'll save actual SELinux configuration for subsequent columns. But, if you have an urgent need to get something working on an SELinux-enabled system, see the on-line Resources for this article.

The Problem

So, precisely what problem are we trying to solve with SELinux? Nothing less than the entire security-patch rat race!

As I've said previously in this space, Linux security often seems to boil down to a cycle of researchers and attackers discovering new security vulnerabilities in Linux applications and kernels; vendors and developers scrambling to release patches, with attackers wreaking havoc against unpatched systems in the meantime; and hapless system administrators finally applying that week's or month's patches, only to repeat the entire trail of tears soon afterward. This is the security-patch rat race, and it's unwinnable. There will always be zero-day (as-yet-unpatched) vulnerabilities.

That's why I've spent so much ink over the years extolling techniques such as virtualizing servers, creating chroot jails, running processes as unprivileged users and using mandatory access controls, all of which limit the effects of zero-day vulnerabilities. SELinux, like Novell AppArmor, is a mandatory access control implementation that doesn't prevent zero-day attacks, but it's specifically designed to contain their effects.

Why is the patch rat race unwinnable? Because in Linux's default Discretionary Access Control (DAC) model, each process runs with the privileges of whichever user starts (or, sometimes, owns) it—that is, *all* of that user's privileges. If an attacker compromises any process running as root, or escalates a compromised process to root privileges, the attacker can do anything root can do, even when that action has nothing whatsoever to do with the process' intended function.

For example, suppose I have a daemon called blinkend that is running as the user someguy, and this daemon is hijacked by an attacker. blinkend's sole function is to make a keyboard LED blink out jokes in Morse code, so you might think, well, the worst the attacker can do is blink some sort of insult, right? Wrong. The attacker can do anything the someguy account

can do, which might include everything from executing the Bash shell to mounting CD-ROMs.

Under SELinux, however, the blinkend process would run in a narrowly defined domain of activity that would allow it to do its job (blinking the LED, possibly reading jokes from a particular text file, and so forth). In other words, blinkend's privileges would not be determined based on its user/owner; rather, they would be determined by much more narrow criteria. Provided blinkend's domain was sufficiently strictly defined, even a successful attack against the blinkend process would, at worst, result in naughty Morse-code blinking.

That, in a nutshell, is the problem SELinux was designed to solve.

What SELinux Does

I'm going to assume you understand how Discretionary Access Controls, aka plain-old filesystem permissions, work in Linux. If you don't, I covered this topic in the October and November 2004 issues of *Linux Journal*, in the two-part series "Linux Filesystem Basics" (see Resources).

Suffice it to say that even under SELinux, the Linux DACs still apply. If the ordinary Linux permissions on a given file block a particular action (for example, user A attempting to write file B), that action still will be blocked, and SELinux won't bother evaluating that action. But, if the ordinary Linux permissions allow the action, SELinux will evaluate the action against its own security policies before allowing it to occur.


So, how does SELinux do this? The starting point for SELinux seems similar to the DAC paradigm: it evaluates actions attempted by subjects against objects.

In SELinux, subjects are always processes. This may seem counterintuitive. Aren't subjects sometimes end users? Not exactly—users execute commands (processes). SELinux naturally pays close attention to who or what executes a given process, but the process itself, not the human being who executed it, is considered to be the subject.

In SELinux, we call actions permissions, just like we do in the Linux DAC. The objects that get acted on, however, are different. Whereas in the Linux DAC model, objects always are files or directories, in SELinux, objects include not only files and directories but also other processes and various system resources in both kernel space and user land.

SELinux differentiates between a wide variety of object classes (categories)—dozens, in fact. You can read the complete list on the Web site "An Overview of Object Classes and Permissions" (see Resources). Not surprisingly, file is the most commonly used object class. Other important object classes include the following:

- dir
- socket
- tcp_socket
- unix_stream_socket



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

AoE is your answer!

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

www.coraid.com



EtherDrive® SR1520

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



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

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



1.706.548.7200

The Linux Storage People

www.coraid.com

- filesystem
- node
- xserver
- cursor

Each object class has a particular set of possible permissions (actions). This makes sense. There are things you can do to directories, for example, that simply don't apply to, say, X servers. Each object class may have both inherited permissions that are common to other classes (for example, read), plus unique permissions that apply only to it. Just a few of the unique permissions associated with the dir class are as follows:

- search
- rmdir
- getattr
- remove_name
- reparent

Don't be frustrated by my not explaining these class names or actions; at this point you don't need to understand them for their own sake. I'm simply illustrating that SELinux goes much, much further than Linux DAC's simple model of users, groups, files, directories and read/write/execute permissions.

As you might guess, SELinux would be impossible to use if you had to create an individual rule for every possible action by every possible subject against every possible object. SELinux gets around this in two ways: 1) by taking the stance "that which is not expressly permitted, is denied" and 2) by grouping subjects, permissions and objects in various ways. Both of these points have positive and negative ramifications.

The "default deny" stance allows you to have to create rules/policies that describe only the behaviors you expect and want, instead of all possible behaviors. It's also, by far, the most secure design principle any access control technology can have. However, it also requires you to anticipate all possible allowable behavior by (and interaction between) every daemon and command on your system.

This is why the "targeted" SELinux policy in Red Hat Enterprise Linux 4 and Fedora Core 3 actually implements what amounts to a "restrict only these particular services" policy, giving free rein to all processes not explicitly covered in the policy. No, this is not the most secure way to use SELinux, and it's not even the way SELinux was originally designed to be used. But as we'll see, it's a justifiable compromise on general-purpose systems.

The upside of SELinux's various groupings (roles, types/domains, contexts and so on) is, obviously, improved efficiency over always having to specify individual subjects, permissions and objects. The downside is still more terminology and layers of abstraction. Alas, with power comes complexity.

So, how does SELinux group subjects, permissions and objects?

Security Contexts: Users, Roles and Domains

Every individual subject and object controlled by SELinux is governed by a security context, each consisting of a user, a

role and a domain (also called a type).

A user is what you'd expect: an individual user, whether human or daemon. However, SELinux maintains its own list of users separate from the Linux DAC system. In security contexts for subjects, the user label indicates which SELinux user account's privileges the subject (which, again, must be a process) is running. In security contexts for objects, the user label indicates which SELinux user account owns the object.

A role is sort of like a group in the Linux DAC system, in that a role may be assumed by any of a number of pre-authorized users, each of whom may be authorized to assume different roles at different times. The difference is that in SELinux, a user may assume only one role at a time, and may switch roles only if and when authorized to do so. The role specified in a security context indicates which role the specified user is operating within for that particular context.

Finally, a domain is sort of like a sandbox: a combination of subjects and objects that may interact with each other. Domains are also called types, and although domains and types are two different things in the Flask security model (on which the NSA based SELinux), in SELinux domain and type are synonymous.

This model, in which each process (subject) is assigned to a domain, wherein only certain operations are permitted, is called Type Enforcement (TE), and it's the heart of SELinux. Type Enforcement also constitutes the bulk of the SELinux implementation in Fedora and Red Hat Enterprise Linux.

There's a bit more to it than that, but before I go any further, I want to use an example scenario to illustrate security contexts.

Suppose we're securing my LED-blinking daemon, blinkend, with SELinux. As you'll recall, it's run with the privileges of the account someguy, and it reads the messages it blinks from a text file, which we'll call /home/someguy/messages.txt.

Under SELinux, we'll need an SELinux user called someguy (remember, this is in addition to the underlying Linux DAC's someguy account—that is, the one in /etc/passwd). We'll also need a role for someguy to assume in this context; we could call it blink_r (by convention, SELinux role names end with _r).

The heart of blinkend's security context will be its domain, which we'll call blinkend_t (by convention, SELinux domain names end with _t—t is short for type). blinkend_t will specify rules that allow the blinkend process to read the file /home/someguy/messages.txt and then write data to, say, /dev/numlockled.

The file /home/someguy/messages.txt and the special file /dev/numlockled will need security contexts of their own. Both of these contexts can probably use the blinkend_t domain, but because they describe objects, not subjects, they'll specify the catch-all role object_r. Objects, which by definition are passive in nature (stuff gets done to them, not the other way around), generally don't assume meaningful roles, but every security context must include a role.

Decision Making in SELinux

There are two types of decisions SELinux must make concerning subjects, domains and objects: access decisions and transition decisions. Access decisions involve subjects doing things to objects that already exist or creating new things that remain in the expected domain. Access decisions are easy to understand. In our example, "can blinkend read /home/someguy/messages.txt?" is just such a decision.

Transition decisions, however, are a bit more subtle. They involve the invocation of processes in different domains than the one in which the subject process is running or the creation

of objects in different types than their parent directories. (Note: even though domain and type are synonymous in SELinux, by convention we usually use domain when talking about processes and type when discussing files.)

That is to say, normally, if one process executes another, the second process will, by default, run within the same SELinux domain. If, for example, blinkend spawns a child process, the child process will run in the blinkend_t domain, the same as its parent. If, however, blinkend tries to spawn a process into some other domain, SELinux will need to make a domain transition decision to determine whether to allow this. Like everything else, transitions must be authorized explicitly in the SELinux policy. This is an important check against privilege-escalation attacks.

File transitions work in a similar way. If a subject creates a file in some directory (and if this file creation is allowed in the subject's domain), the new file normally will inherit the security context (user, role and domain) of the parent directory. For example, if blinkend's security context allows it to write a new file in /home/someguy/, say, /home/someguy/error.log, then error.log will inherit the security context (user, role and type) of /home/someguy/. If, for some reason, blinkend tries to label error.log with a different security context, SELinux will need to make a type transition decision.

Get the picture? Transition decisions are necessary because the same file or resource may be used in multiple domains/types; process and file transitions are a normal part of system operation. But, if domains can be changed arbitrarily, attackers will have a much easier time doing mischief.

Conclusion

Besides Type Enforcement, SELinux includes a second model, called Role-Based Access Control (RBAC). Although I'm out of space for now, RBAC builds on the concepts we've already discussed, providing controls especially useful when real human users, as opposed to daemons and other automated processes, are concerned.

Next time, I'll describe RBAC at length and begin going into greater depth on how actually to use SELinux, beginning with Fedora and Red Hat's "targeted" policy. Until then, be safe! ■

Resources for this article: www.linuxjournal.com/article/9510.

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".



Expert Included.

Jonathan provides expert, dedicated technical support for one of the most comprehensive server and storage product offerings in the industry. He is a fan of the Rackform iServ R276-SAS storage server with two Quad-Core Intel® Xeon® Processors 5300 Series because it combines Intel's proven reliability with industry-leading 8-core performance. And flexible, fault-tolerant data protection ensures that your server stays up and running, so you don't have to call Jonathan.

When you partner with Silicon Mechanics, you get more than a super Intel solution — you get an expert like Jonathan.



Quad-core. Unmatched.



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

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



JON "MADDOG" HALL

Beneath the Surface

Don't forget the modular command-line power of Linux.

I was walking along the beach with one of the Pollywogs when I saw a small tidal pool. I stopped to wade through it and look at some of the life under the rocks.

Most people never look under the rocks in a tidal pool or in a freshwater stream, but there is a lot of very interesting and necessary life to be found—life forms that are necessary because they fill a very important part of the world. Most people see only the glossy surface of the ocean or the stream, simply because they never look any deeper.

The same is true with Linux. I have noticed that recently there has been a lot of work on graphical user interfaces, with translucent windows and different ways of displaying multiple desktops—all of this is good.

In my opinion, however, the real power of Linux comes from the command-line interface that resides below this glossy surface and allows people to write very powerful programs to manipulate huge amounts of data.

I do not expect that everyone will want to learn every type of command-line interface or small language, but if you do not learn at least one or two, you will never know how powerful your system can be.

In my opinion, however, the real power of Linux comes from the command-line interface that resides below this glossy surface and allows people to write very powerful programs to manipulate huge amounts of data.

Many years ago, the company where I was working needed to get a new piece of software out to its customers. However, the customers who were supposed to receive the software were represented by two different printouts from two different systems, and my company was planning on having a clerk evaluate the two reports to accomplish this task. Estimated time for the clerk to do this was nine months, which meant that the software would be almost a year old before the customers received it.

I asked if this process could somehow be automated, because the customers were waiting for the software. "No", I was told, "it can't be done", because the databases were incompatible and on different machines. There was no program that could reach across the systems to coordinate the data.

I had the managers put the printout into two files, and put both files on my (at that time) UNIX system. In less than a quarter of a day, using the stream editor `sed(1)`, the pattern matching program `grep(1)` and the pattern matching, scanning and processing language `awk(1)`, I was able not only to correlate the data but also to print out mailing labels for the shipping boxes along with an indication of the proper software to go in each one. The managers could not believe it.

Some people think that it takes a lot of study in order to "know" command-line programming. However, if you approach the task systematically, you can learn it over time, taking advantage of each learning cycle.

The first thing you probably should do is get a book on Linux commands. *Linux In A Nutshell: A Desktop Quick Reference* by Figgins, Weber and Siever (O'Reilly) is a good start. Another good one is *Linux Pocket Guide* by Barrett, also from O'Reilly. Finally, *Linux For Dummies Quick Reference* by Hughes and Navratilova (Wiley) also is a good reference.

Read the book you choose, but do not obsess with memorizing the capabilities of each command. After you have read the book, think about some task you have to do repeatedly and what it would take to automate that task. You probably will find some Linux command-line programs that would help make things easier.

When you log in to your Linux system, execute a terminal emulator program, such as `xterm` or one of the others. Stay away from superuser (`root`) mode for the present, as you are trying to learn and sometimes things go astray.

Practice with some commands, such as `grep`, `sed`, `ls`, `cd` and others, simply by typing them into the command line and feeding them data according to what the command requires. Or, create a file of ASCII characters that you would like to use the commands to search, sort, filter or otherwise change.

Then, start putting the commands together using the pipe symbol (`|`). Note that this is not either the lowercase `l` or uppercase `i`. It is typically found along with some of the other special characters on ASCII keyboards, usually above the Enter key.

For example, start by putting together the `ls` and `grep` commands:

```
ls | grep 'e'
```

This will show you every visible file in your directory with the letter `e` in its name.

Another area of study should be the concept of regular expressions—ways of describing strings of data that typically are used for searching or matching with other strings of characters. The aforementioned books also cover issues of regular expression creation, which can be quite tricky, but also quite powerful.

Although different programs may use different methods of regular expressions, they tend to follow the same principles, and generally you can use the same type of special characters with each command.

I was working for Bell Laboratories in 1977, trying to be a system administrator for this interesting system called "UNIX". For several months I had been frustrated by trying to learn this operating system that had seemingly millions of tiny little commands, multiple directories holding them and "cryptic" names for them. One night I was trying to modify a text file with the interactive text editor, `ed(1)`, and I could see that it would take me *hours* to modify the file using `ed`, if not all night.

I remember suddenly thinking, "I do not *know* that there is a command in UNIX for doing this easily, but I am willing to

bet there is one." So, I started going through the manual looking only at the description of each command given in the "Name" line for the command. Fairly soon, I came across cut and its partner program paste, which allowed me to do exactly what I needed to do in two commands. From that time on, I followed the philosophy of first looking for the right command, and although that philosophy was sometimes wrong, more times than not, the philosophy was right, and a suitable command did exist.

To start learning the command line with only on-line resources, make sure that you have loaded the on-line manual and info pages from your distribution. You can then type in `man intro` to read the introduction section of the `man(1)` command, then type `man <command-name>`—for example, `man ls`—to learn more about the `ls(1)` command. The (1) after the command name `ls` means that it is a user-level command, rather than a programming interface, system administrator command or other specialized function.

If you like a graphical, mouse-based reader, rather than a command-line reader, there is `xman`. Once you have invoked `xman` by typing `xman`, click Help in the little window and read the first section of the help page. You then can click manual page in the little control window, and when the text window

pops up, select show both screens from the Options menu at the top. This lets you see both the index of all the manual commands in the top section and the actual manual page itself in the bottom section. Click on the program of interest in the top section, and the command will be formatted in the bottom section. An example of an interesting command is `less(1)`.

I can't touch on all the issues and needs for learning the power of the command line in one column, but perhaps I've piqued your interest in discovering why a lot of Linux users do not use a graphical windowing system at all, preferring to use only the command line, while others (myself included) heavily use both the windowing system and the command line.

And, perhaps you will look beneath the surface to see the power of the underlying currents. ■

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



Don't complicate a simple task

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



M7100 Wireless Family

AML
NEVER COMPLICATED™





DOC SEARLS

Migrating a Mentality

We're not going to get the Net we want until we quit thinking it's gravy on top of telephone and cable service.

The Internet will do for the 21st century what roads did for the 20th century and railroads did for the 19th century. That we need to build out the Net, to the maximum possible capacity, everywhere we can, is beyond question. That economic and cultural benefits will increase with connectivity and capacity is also beyond question. What's not beyond question is who should do it, how, where and by when.

In Korea, Japan, Denmark, the Netherlands and other countries, there is widespread public and private commitment to build out Net connectivity to as many people and places as possible, with as much capacity as possible. Means differ, but the goals are the same. Net build-out is a top priority.

Meanwhile, here in the US, Net build-out has been left up to cable TV and telephone companies that have not only squandered opportunities (according to TeleTruth, carriers have pocketed \$200 million in federal subsidies for fiber build-outs that never happened), but have conflicted interests in the matter. Here's how home networking pioneer (and co-inventor of the spreadsheet) Bob Frankston puts it:

For those worried about competition, it would be hard to do worse than a system in which there is a fundamental conflict of interest. Today's transport providers have a very strong incentive, even a requirement, to maintain scarcity—especially when burdened with costs that do not increase the value of their product.

This is why fiber deployments like Verizon's FiOS are really about delivering high-definition television (and competing with cable TV companies), rather than delivering Internet capacity. Bob continues:

The fiber they are installing for FiOS is really a cable TV plant disguised as a network. It is a Passive Optical Network (PON) designed as a distribution system from a head end to the terminals at each home; though it does have capacity to send data back. A single fiber has the capacity for gigabits of traffic. There's so much capacity that they can simply allocate a portion of the capacity to emulating traditional cable TV. The 15mbps they reserve for their Internet service is less than 1% of that capacity!

He adds, "Direct and transparent funding is vital, but unlike the current regulated system, we do not have to grant the transport providers any exclusive rights—we can all add capacity."

This is the key point. Adding infrastructural capacity for Internet isn't as hard or complex as building roads, bridges, dams, waste treatment facilities, railroad lines or power plants with large towers marching across the landscape. It's mostly a matter of planting conduit and fiber-optic cabling in the ground, or hanging cabling from poles that are already there—then deploying wireless coverage with fiber "backhaul".

Bob Frankston's preference is for individuals and communi-

ties to build their own DIY (do-it-yourself) "plant" and connect in their own ways with each other, bypassing the cableco/telco duopoly and the "regulatorium" (his word, and it's an excellent one) that governs it. Local DIY networking is exactly the business of Indinet.dk in Copenhagen, which I wrote about last month. Not surprisingly, the Organisation for Economic Co-operation and Development (OECD) lists Denmark as the top country in broadband penetration and growth. The US is 12th in penetration and 17th in growth.

Here in the US, citizens are opting to use local governments for DIY Net build-out. The results are "muni" projects by cities and counties—hundreds, so far, across the country. In New Mexico, Sandoval County—home to seven Native American pueblos, a 33% Hispanic population and Intel's largest fabrication plant—is spending around \$8 million (a remarkably low number) on a wireless build-out that intends to deliver gigabit-level connectivity to everybody in a region the size of Connecticut, yet notoriously lacking in amenities. In Utah, UTOPIA is a fiber build-out by 14 cities that wholesale capacity to retail service providers. One of those, ironically, is AT&T. In Vermont, Burlington Telecom is a city department currently building out a "triple play" (Internet, phone, television) retail offering.

Each project is unique, but all have two things in common: 1) they're doing what the carriers won't, and 2) they're doing it for every citizen, organization and business—and not for one company or one application.

Naturally, the carriers oppose the munis. They say these local governments are competing with business (which is highly ironic, given that the carriers have lived under government-maintained regulatory protection for the duration). So the carriers have been lobbying for anti-muni legislation at the state and federal levels. One of their successes is the Local Government Fair Competition Act in Louisiana, which was passed at the behest of the carriers to "level the playing field" between them and the munis. The law has had the effect, so far, of halting deployment of a fiber-based muni system in Lafayette that originated with voters.

Everywhere you look, the carriers are at odds with their own customers. Last November, 72% of the voters in Clarksville, Tennessee, approved the city's Department of Electricity's bid to build out a fiber-based network.

The arguments are not going to get any less heated, especially with a new US Congress that features a Democratic Party majority. In the last Congress, Net Neutrality legislative efforts, led by Democrats, were defeated by Republican majorities. Pro-Neutrality advocates will be looking for new legislation to be introduced. And, you can bet the carriers will fight that legislation by stepping up PR as well as lobbying efforts.

We can get past those arguments and simplify matters by answering one deep and simple question: is the Net public or private infrastructure? The munis say public. The carriers say private. To help find the answer, here is a list of familiar infrastructures, sorted into public, private and a mix of both.

Do you take

"the computer doesn't do that"

as a personal challenge?

So do we.

LINUX
JOURNAL™

Since 1994: The Original Monthly Magazine of the Linux Community

Subscribe today at www.linuxjournal.com

Public:

- Water (wells, reservoirs, distribution systems, dikes and levees).
- Streets, roads, highways and bridges.
- Waste water treatment.
- Garbage disposal (mostly landfills).

Mixed public and private:

- Garbage collection and recycling.
- Electric power generation and distribution.


Private:

- Telephony.
- Broadcasting.
- Cable TV.

We can argue about what belongs on the list and what doesn't. But what's clear is that we need public infrastructures to support civilization.

Public infrastructure is manufactured nature. Reservoirs are man-made lakes. Irrigation canals are man-made streams. Waste treatment systems are man-made swamps. Roads and bridges are man-made geology. Power-generating plants are man-made systems for converting or extracting energy from nature. At their best, public infrastructures work as part of nature. Water capture, distribution and waste treatment should work inside the hydrologic cycle. Roads and bridges should conform to the supportive shapes and materials that make up the world's lands and waters.

You can make money with public infrastructure, but that's not infrastructure's main purpose. What you want is to make money because of infrastructure. Roads, water and waste treatment are all built primarily to support economies other than their own, if they even have any. Even our electric and gas utilities are not in business to support only themselves. They are in business because the rest of civilization can't get along without them. Public infrastructures are so quietly supportive to civilization that most of us give no more thought to them than we give to gravity or sunlight.



The Net is quietly supportive too. It doesn't advertise itself. It only connects devices and carries bits. It reduces to zero the distance between any two devices, or any two individuals. What we get billed for by phone and cable companies is access to the Net—not the Net itself.

I would argue that the Net is the most public infrastructure we've ever built, because it's the first to build on human nature. To illustrate this, Figure 1 is a diagram of civilization, borrowed from the Long Now Foundation.

I've shown this before, but I think it's important to show it again, because it shows how each layer supports the one above it, allowing the higher layer to move faster.

Let's look at the case of Linux, which grew out of the need to develop tools and building materials that are useful to everybody, rather than to just one company. This universality of purpose is what makes Linux infrastructural. The natural way Linux (and other open-source tools and building materials) grows also resembles that of a species. Here's how I explained this in a report last year:

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

The species-like nature of FOSS (free and open-source software) is organized by community development culture, which gives rise to self-governance within communities—along with licensing that makes infrastructural choices as solid and useful as possible to commerce, to markets, to entire economies. Thus, infrastructure arises out of, and builds upon, the best of human nature.

All this was clearly evident last November, when I walked around the exhibition hall at ISPCON. Dozens of infrastructure deployment businesses (mostly selling local and regional wireless Internet equipment) built their systems on Linux. When some of folks at the booths saw *Linux Journal* on my badge, they wanted to tell me how they put Linux to use. In other cases, I had to ask. Usually the answer was "Oh, sure." It was like asking if they wore clothes. The answer was that obvious.

Linux became ubiquitous because experts put it to use. Experts discovered the benefits quickly, and expertise around Linux eventually became a premium skill set. Jakob Frederiksen of Indinet.dk told me that Linux talent was cheap five years ago, but expensive today. (This is one more example of making money because of Linux rather than just with it.)

"All the significant trends start with technologists", Mark Andreessen told me 11 years ago (when Netscape open-sourced Mozilla). He also said, "Technologists are driving progress, and it's easier to drive with Linux than with anything else."

There is a lag between what technologists do first, and what the rest of us do later—especially when what technologists do is not strictly commercial, yet is deeply supportive of commercial activity. The way nature, culture, governance and infrastructure all support commerce is not apparent at the commercial level. Nor is the way commerce contributes back to infrastructure as well. Yet we can be sure that the experience of many Internet infrastructure builders in the world will contribute useful code to Linux and many other infrastructural building materials and tools.

Meanwhile, most business experts still don't grok the infrastructural nature of the Net, even though they put it to

Figure 1. Diagram of Civilization (from the Long Now Foundation)

use every day. Like most of the rest of us, they're still stuck in the Net's equivalent of the 1880s, when electric power was just beginning to replace gas, and most people understood electricity in terms of its primary use, which was light. Even today, many electric utilities still carry the surname "Power & Light". DC vs. AC was the Cable vs. Telco of its day.

In the long run, we learned to separate power from light—or, in modern parlance, transport from applications. As Bob Frankston puts it, "Edison originally sold light, but we now buy electricity and create our own lighting." Today the equivalent of "light" for most of us is a combination of e-mail and Web browsing. A guy selling business-grade Internet service for our local cable company (Cox Communications) told me recently that most new business Internet customers use the Net to connect retail point-of-sale devices, and to do a combination of e-mail and browsing in their offices. They haven't discovered the full potential of high-speed symmetrical Internet service.

Of course, the carriers have hardly given any of us the chance. They have ignored the fact that the Net was designed in the first place as a symmetrical system, with equally fast and unencumbered upstream and downstream connections and speeds. As a result, almost none of us with a home or low-end business connection has ever experienced symmetrical service. The carriers optimized their systems from the beginning to anticipate and support consumption, not production. Moreover, business customers were charged a premium, just like they've always been charged premiums for "business" telephone and cable TV service.

Now let's talk about cost.

Fiber isn't free, but it's generally cheaper than the cost of planting it underground or hanging it from poles, and it's getting cheaper every day. More important, each strand can carry gigabits of data. The "first cost" of the Net, once fiber is installed, is blinking light. Routers, amplifiers and other infrastructural gear cost money to buy and to run, but the costs of the connections themselves are basically zero. And fiber cabling doesn't deteriorate with use. That's because there is no physical difference between fiber that's "dark" and fiber that's "lit". Light does dim over distance, but it doesn't encounter the same degrading resistance that electrons meet as they pass through copper wiring. Fiber optic signals also emit no side radiation along the cabling. So, it's also about as "green" as a technology can get.

Wireless deployments are cheaper than fiber (no need to trench or hang cabling), and are capable of spanning distances where fiber deployments are impractical or impossible (such as across canyons of the Southwest US). But in both cases, the investments are highly durable, far less costly than most highway, water and waste treatment projects, and hugely supportive of countless activities, and markets of every sort.

The top price for FTTx (fiber to the whatever) that I've heard so far is about \$2,500 per "drop". This is about the same price you'll pay for a big flat TV screen that will be obsolete in three years, if not less. Meanwhile FTTx will only improve in value.

What about funding? Bob Frankston says, "Financing fungible connectivity in the same way you might finance macadam makes sense. Financing streets based on being able to stop cars and demand protection money is very different."

The problem with "triple play" for munis is that it puts them in direct competition with their local telephone and cable companies. Worse, it makes communities come up with a commercial "revenue model" for public infrastructure. We don't burden roads and water systems with revenue models that do anything more than cover the expenses of maintaining them. Why should we place that burden on the Net?

Because the only models we know are provided by phone and cable companies. Also because we want to pay off debts, and "triple play" seems like a good way. Unfortunately, by emulating the carriers we not only adopt their business models but also their mentality. "Triple play" sees only three ways of making business with the Net, rather than limitless ways of making money because of the Net. By building out the Net, we're creating an ocean of connectivity, with frontage for everybody. The ocean's job is to support every kind of use, every kind of traffic, every application, every business, equally.

Perhaps the best model for munis is the municipal electric utility. Jim Baller, one of the top lawyers specializing in muni build-outs, writes:

More than 2,000 municipal electric utilities have thrived over the last century, contributing greatly to the well-being of their communities and America as a whole. Another 1,000 communities established their own electric utilities and sold them to the private sector, having achieved their goal of avoiding being left behind in obtaining the benefits of electricity. In contrast to these 3,000 successful municipalities, thousands of communities that waited for the private sector to get around to them stagnated or became ghost towns.

We are at the same crossroads today—except that only one road is built, and we need to build the other road across it. ■

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



Linux Certified

Linux Laptops: The New LC2000 Series

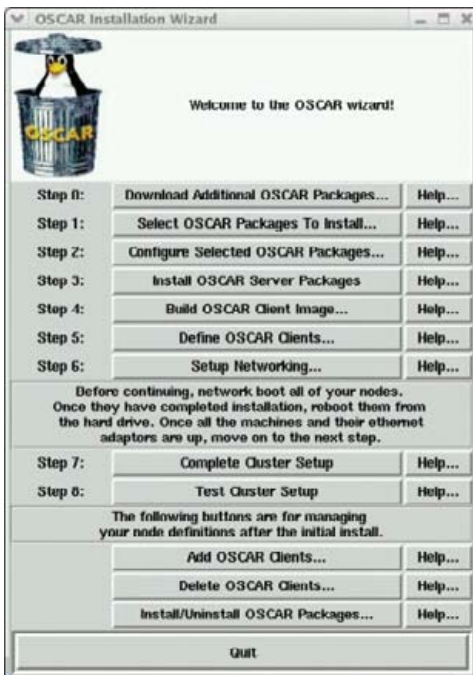
- *High Performance*
- *Amazing ROI*
- *Robust*
- *Fully Compatible*
- *Cost Effective*

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

OSCAR Working Group's Open Source Cluster Application Resources Toolkit

Cluster gurus rejoice! The OSCAR working group recently released version 5.0 of its Open Source Cluster Application Resources (OSCAR) toolkit. OSCAR is a software package that "supports the use of high-performance computing by reducing the work of cluster configuration, installation, operation and management." OSCAR's developers have revamped the application's infrastructure and included many new features, such as smart package managers, yum-based package installs and image building, easier client updating via a new repository approach and optimized startups to reduce build times. A new utility called netbootmgr has been added, which "greatly reduces the amount of time spent mucking about in the BIOS by centrally managing a nodes behavior when a network boot is detected". In preparation for future releases, a new package and database structure has been designed in anticipation of Debian support. OSCAR 5.0 also has been fully tested for use with both IA32 processors and x86_64 processors under several major Linux distros. OSCAR is available for download from the group's Web site.

oscar.openclustergroup.org

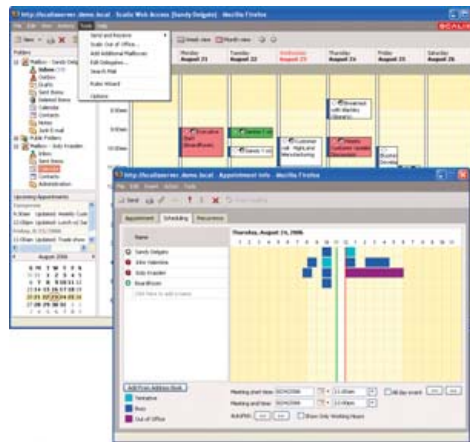


Message Partners' MPP

Got Postfix? Then Message Partners wants you to use its MPPv3, the company's integrated pre- and post-queue spam filter for Postfix. Message Partners claims that MPP

"solves all of the problems that complex e-mail environments run into in a single high-performance application", including solutions for virus and spam filtering, content filtering, access controls, end-user quarantine and white/black list management, archiving and other features. A key new feature is the Postfix Policy Server, which adds the "capability to make pre-queue admission decisions for every type of e-mail (including multirecipient and multidomain)." Message Partners also touts its innovative sharing of a common database and configuration by the pre-queue and post-queue filters, which improves management of per-domain SMTP restrictions in large environments. A free trial is available at the company's Web site.

www.messagepartners.com



Scalix

Our contacts at Scalix Corporation informed us about Scalix 11, the company's Linux-based, open-source supported messaging, e-mail and calendaring platform. Some of Scalix's main features include easy administration, "deep integration with legacy environments" and Outlook-level functionality without the costs and license lock-in of MS Exchange. Scalix's target customers are those requiring the "product integrity of an enterprise platform with the community support of an open-source project". New features include two new Web services, a lightweight mobile client, enhanced management capabilities, and improved Web client and Outlook support. Vis à vis the new Microsoft Exchange Server 2007, Scalix claims advantages, such as "vendor choice, better administration and broader client selection while maintaining the best Outlook support in the market". A trial of the commercial edition and a free community edition are available for download from Scalix's Web site.


www.scalix.com



VMLogix's LabManager

Virtualization is dynamic these days, and VMLogix adds to the fabulous ferment with LabManager, the company's virtualization solution centered on the software development life cycle. LabManager, says VMLogix, offers "rapid, highly repeatable, resource-optimized deployments of complex, multemachine software test environments" that allow developers to reduce cycle times and infrastructure costs, share resources seamlessly and improve the quality of the final product. LabManager's approach enables users to exchange development environments with each other regardless of time, location or project status. Other product advantages include platform agnosticism, physical- and virtual-machine capabilities and compatibility with most testing and development tools. Free evaluation licenses and downloads of LabManager are available at VMLogix's Web site.

www.vmlogix.com



Pete Goodliffe's Code Craft (No Starch Press)

It's always a treat to buzz over to No Starch and see the latest goodies for geeks. If you visit No Starch now, you'll find Peter Goodliffe's brand-new book *Code Craft: The Practice of Writing Excellent Code*. The book's purpose is to take the programmer to a new level, from writing correct code to writing great code that is easy to understand. *Code Craft* is language-agnostic and covers not only issues such as presentation style, variable naming, error handling and security, but also effective teamwork, development processes and documentation. A free sample chapter is available at No Starch Press' Web site.

www.nostarch.com

Terra Soft's Yellow Dog Linux for PlayStation 3

Linux people simply rock, don't they? A case in point is Terra Soft, which has ported its Yellow Dog Linux (YDL) v5.0 to the PlayStation 3 (PS3) from Sony. Hopefully by the time you read this, you can finally get your hands on a PS3, because we can't! "With Sony's blessing and support from the Barcelona Supercomputing Center, YDL is a full Linux OS for PS3, which is based on Fedora Core 5 and comes complete with more than 1,500 packages. In this project, Terra Soft also collaborated with Carsten Haitzler and the Enlightenment development team to integrate the E17 desktop, which the firm says will provide "an unprecedented level of function and interface aesthetic". Install and source ISOs for YDL are available for download; DVDs are available for purchase from the Terra Soft's on-line store.

www.terrasoftsolutions.com



ADLINK Technology's NuPRO-851

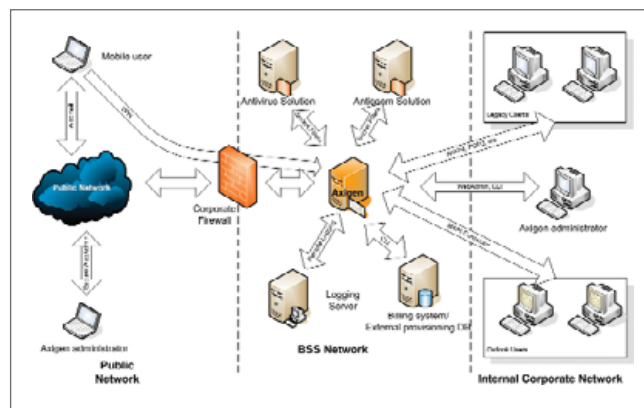
ADLINK Technology just brought forth a new full-sized, single-board computer, the NuPRO-851 Series. The NuPRO-851 Series is a PICMG 1.0 device that supports 800MHz FSB with a hyper-threading Intel Pentium 4 (LGA775) processor, dual-channel DDR2 memory at speeds of 400/533MHz and an Intel GMA900 graphics core architecture providing up to 2048x1536 resolution with 8.5GB/s peak memory bandwidth. The product also features the Intel 915GV and ICH6 chipsets, USB 2.0 connectivity and two on-board Marvell 88E8052 controllers that support dual-gigabit Ethernet ports via the PCIe bus. ADLINK says the product is "ideal for industrial controllers and equipment providers", because it places heavy emphasis on longevity, reliability and strict revision control. This and all other ADLINK products comply with the European Union's RoHS directive on environmentally sound products.

www.adlinktech.com

AXIGEN's Mail Server

Just when you thought that the dizzying array of options for mail servers couldn't get more dazzling, AXIGEN releases version 2.0 of its Mail Server product. The package provides "all-in-one server functionalities, ranging from e-mail communication to anti-virus and anti-spam integration." In the new edition, AXIGEN has added many features, including a backup and restore module, a full-security toolbox that integrates 16 anti-virus and anti-spam applications, localized and skinnable Web mail, a reporting engine for more than 100 definable reports, a wizard for creating e-mail delivery rules and others. The Mail Server is currently available in both the English and German languages and runs on major Linux distributions. An evaluation edition is available at AXIGEN's Web site.

www.axigen.com



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

THE INTEROPERABILITY POWER OF LINUX-NTFS TOOLS

SOME OF THE BEST WINDOWS UTILITIES
AVAILABLE MIGHT BE IN YOUR LINUX PARTITION.

Steven Mathes

A Linux environment can gain access to dozens of filesystems, whether on the local hard drive or somewhere on the network. More specifically, Linux can run many tools to manipulate Windows filesystems or repair Windows problems.

One suite of tools comes from the Linux-NTFS Project. These utilities work many miracles. One resizes NTFS partitions. Several manipulate individual files. One clones an entire NTFS image. It is possible to back up Windows installations, clone new workstations from a centrally stored image and update images across a network. And, because these tools run inside Linux, they benefit from the power of the Linux environment. These tools help when you're dealing with a single dual-boot computer. They quickly become indispensable if you work with a large network. Aided by redirection, pipes and scripting, it is easy to automate many tedious but important Windows maintenance tasks from within Linux.

Installation

The utilities are widely available and well supported. Packages are available for virtually all Linux distributions that have package managers, and the software itself is even included on the Knoppix live CD. Many distributions install the tools to be run only by the root user. To see if these tools are on your installation of Linux, consulting the man pages will at least show whether the documentation is installed: `man ntfsprogs`.

Even if the software and/or documentation are absent, you can install these tools yourself. For SUSE, Debian, Ubuntu and Gentoo, `ntfsprogs` is the package name to search for and install. The packages for some distributions include all of the NTFS tools, some do not. For example, the package in the Etch version of Debian includes the `ntfsmount` tool, and the package in the Sarge version does not. Red Hat/Fedora distributions do not support NTFS, based on perceived licensing issues, but specifically designed packages for Red Hat/Fedora are available directly from the Linux-NTFS Project. Of course, consulting the actual home page of the project (www.linux-ntfs.org) gives the most up-to-date documentation and information, as well as the latest source code and instructions for building the complete set of tools.

No matter what flavor of Linux you run, it is possible to download the source code and install from that. This is a good choice if you want the newest features and the latest NTFS drivers, although you could suffer from the disadvantage of having bypassed your package manager.

Note: before you build `ntfsprogs` from scratch, you probably should install the FUSE library (fuse.sourceforge.net). Linux has a built-in NTFS driver, but the NTFS utilities include a second driver for NT filesystems. The non-native driver is the FUSE-based `ntfsmount`, which boasts many extra features. However, it is a bit slower than the driver that comes with the latest kernel. Furthermore, it requires that your kernel has the FUSE module.

If you want to install the FUSE library, download the latest source and store it in a handy directory, maybe the same place you plan to store your `ntfsprogs` download. The installation follows the "configure, make, make-install" process that has become the standard (note that the version number may have changed by the time you read this). Do this as root:

```
tar -xzvf fuse-2.5.3.tar.gz
cd fuse-2.5.3.tar.gz
./configure
make
make install
```

Installing the FUSE library and module is not completely necessary if all you want is read access (and somewhat temperamental read/write access) to an NT filesystem. That's because for all distributions, except Red Hat/Fedora, there is a native Linux kernel driver that runs through the normal mount command. It is faster, but it lacks the extensive features and

feedback of `ntfsmount`.

Now, download the `ntfsprogs` source, and then save it in a handy directory. Operating as root, build it much the way you built the FUSE package (again, the actual version number may have changed by the time you read this):

```
tar -xzvf ntfsprogs-1.13.1.tar.gz
cd ntfsprogs-1.13.1
./configure
make
make install
```

When building `ntfsprogs` without the FUSE library (even if you do have the FUSE module), you will get a complaint while running the configure command:

```
checking for FUSE_MODULE... configure: WARNING: \
ntfsmount requires FUSE version >= 2.3.0
```

This shouldn't be fatal to building the other NTFS tools, but you will not be able to compile `ntfsmount`.

If you are running Red Hat/Fedora, you might not even have the kernel driver. In that case, it is strongly recommended that you either install a custom kernel containing the kernel-based NTFS driver or install the FUSE libraries before building.

The Software

At this point, it is assumed that you have either installed `ntfsprogs` or have discovered it already installed on your system.

If you have already looked at the `ntfsprogs` man page, you have seen the complete list of the utilities. Here is that part of the output from the man page:

```
mkntfs(8) - Format a partition using NTFS.
ntfscat(8) - Dump a file's contents to the standard
output.
ntfsclose(8) - Efficiently clone, create, restore or
rescue an image of an NTFS partition.
ntfscluster(8) - Locate the owner of any given sector
or cluster on an NTFS partition.
ntfscp(8) - Overwrite file on an NTFS partition.
ntfsfix(8) - Check and fix some common errors, clear
the LogFile and make Windows perform a thorough check
next time it boots.
ntfsinfo(8) - Show some information about an NTFS
partition or one of the files or directories
within it.
ntfslabel(8) - Show, or set, an NTFS partition's
volume label.
ntfsls(8) - List information about files in a
directory residing on an NTFS partition.
ntfsmount(8) - NTFS module for FUSE.
ntfsresize(8) - Resize an NTFS partition without
losing data.
ntfsundelete(8) - Recover deleted files from an
NTFS partition.
```

Many of the utilities listed are more useful to developers than to people doing maintenance on a network or dual-boot computer. However, some of these are real life-savers, and `ntfsclose` is the biggest life-saver of all.

Using the NTFS Tools

In order to try out `ntfsclose`, you need a computer with the NT filesystem to clone, and it needs to have access to another filesystem large enough to store the image. Recommended filesystems are `ext2`, `ext3`, `xfs` or `ReiserFS`. The documentation for `ntfsclose` warns that the `ReiserFS` is slow when

handling sparse files, but I have found the performance to be okay with more recent versions. It is possible to use an external drive, as long as it has the ability to store huge files—for some operations you will need space as large as your entire Windows partition. If you have an external drive formatted as a FAT32 filesystem, it will have a size limit for individual files that is too small for what you need. Of course, if your computer does not have Linux installed, you will need to boot from a live Linux CD, such as Knoppix.

Notice that the description of the ntfsclone utility above claims that it does its job “efficiently”. This is not merely a boast. On newer hardware, it can clone a substantial Windows XP workstation in just a couple of minutes. If you had an NT filesystem on the first partition of the first IDE drive and were operating from Linux on the same computer, the following command would back up the NTFS as a single file:

```
ntfsclone /dev/hda1 -O /usr/local/backup/ntfs.img
```

The uppercase O in this command tells the software to overwrite the image, but it will create the file if it is absent. This will not compress the filesystem. In fact, it will leave it in a state to allow you to mount ntfs.img using loopback. First, make a mountpoint:

```
mkdir /usr/local/backup/mtpt
```

Then, use ntfsmount and the same syntax you would use for an ordinary mount:

```
ntfsmount -o loop /usr/local/backup/ntfs.img \
/usr/local/backup/mtpt/
```

The ntfsmount command mounts the filesystem read/write by default. Files can be copied, moved and deleted easily. Of course, there are the usual cross-platform perils to contend with. For example, situations involving configuration files can require caution when alien line endings and character sets are involved.

Using the native mount command with the native driver involves the same familiar syntax:

```
mount -o rw,loop,nls=utf8 -t ntfs \
/usr/local/backup/ntfs.img \
/usr/local/backup/mtpt/
```

Note that this mount also makes a provision for a Windows-compatible character set. You still need to use caution, finesse and expertise, however, if you were to choose to edit, say, boot.ini with Emacs. It would be better to edit such a file in a Windows environment or perhaps with Notepad running through Wine.

If you want read/write access, your success with this last mounting method might vary according to the version of your kernel. Again, the native driver is a bit finicky. It may complain, and if it does, its usual behavior is to fall back to a mount that is read-only. Older versions of the native driver are outright dangerous in read/write mode.

Unmount the filesystem the same way for both methods. From the directory containing the mountpoint do the following:

```
umount mtpt/
```

The ntfs.img file can be moved and copied just like any other (admittedly huge) file. It can be compressed and stored in a safe place. It can be uploaded to remote locations. A copy can be edited and then restored over the original. The command for restoring this backup onto the original partition (while in the directory containing the backup) is as follows:

```
ntfsclone ntfs.img -O /dev/hda1
```

Sometimes, smaller is better. The ntfsclone command will take flags that allow your image to be compressed efficiently. These flags also make the process of cloning much faster, both from the local hard drive and over the network. Here is one example, where the image is saved much the way it was in the first example:

```
ntfsclone --save-image /dev/hda1 -O \
/usr/local/backup/ntfs.img
```

This image, alas, cannot be mounted unless it is restored, either to its original partition or to a different file. Restoring to its original partition would happen as follows:

```
ntfsclone --restore-image --overwrite /dev/hda1 \
/usr/local/backup/ntfs.img
```

Note that in the above, the -O has been replaced by the more script-friendly --overwrite flag. They do the same job. All flags can be expressed as script-friendly words (for readers of English), and most can be expressed as single letters.

PEG[®] the leading GUI for Embedded Systems

Prototype Today!! FREE EVALUATION PACKAGE

- PEG+ - Full Featured Windowing in C++
- C/PEG - Smallest Footprint in ANSI C
- Royalty Free
- Fast execution speed
- Completely ROM-able
- Delivered with Full Source Code
- Development Tools including FontCapture, PEG WindowBuilder, and ImageConvert
- Complete set of screen drivers included
- Completely customizable
- Industry leading RTOS Support
- Supports all popular target processors, video controllers and I/O devices
- Multi-lingual support – 2-byte character sets & UNICODE string encoding
- Event-driven programming model
- Application Design Services
- Knowledgeable and timely support to users around the globe
- Now includes a fully licensed version of Paint Shop Pro 9

SWELL software
GRAPHICS SOFTWARE FOR EMBEDDED SYSTEMS
WWW.SWELLSOFTWARE.COM | 810-982-5955

Now comes the good part. The `ntfscdclone` utility will send its data to standard output. This means you have your choice of various compression utilities, different modes of transfer over the network and so forth. Any useful tool that accepts standard input could process the image. Here are some examples.

To back up a compressed image, do:

```
ntfscdclone --save-image --output - /dev/hda1 | gzip \
-c >ntfs.img.gz
```

The image is sent to standard output by the `-output` flag with the argument of a single dash. The `gzip` utility compresses it, then redirects the stream to overwrite or create the file `ntfs.img.gz`.

To back up the image to a remote computer, do:

```
ntfscdclone --save-image -o - /dev/hda1 | ssh \
backups@storage.mydomain.org \
"dd of/home/backups/windows/images/ntfs.img"
```

Here, the flag for `--output` is shortened to its single-letter abbreviation. It is sent to standard output. This, in turn, is piped into the `ssh` program. The stream is sent over the network to a computer named `storage` under the care of a user named `backups` and stored in its proper place through the `dd` command.

Here is another example:

```
wget ftp://storage.mydomain.org/home/backups/
windows/images/ntfs.img.gz \
-O - | gunzip | tee /usr/local/backup/ntfs.img | \
ntfscdclone --restore-image --overwrite /dev/hda1 -
```

This could be a line taken directly from a cloning script, because it needs no password or other user input. It uses `wget` to download the compressed image, uses `gunzip` to unzip it, and then splits the data stream with the `tee` command, so that a backup copy of the image is stored in the Linux partition at the same time that it is redirected to the NT partition on `/dev/hda1`. This assumes that `storage.mydomain.org` has a functioning anonymous FTP daemon. Other possible ways of downloading without user input would be to use `wget` with Apache or to set up encryption keys to use with SSH. Again the possibilities are limited only by the incredible number of tools available.

Another useful tool in the `ntfsprogs` package is `ntfsresize`. This does exactly what it advertises. It shrinks or expands an NT filesystem. It operates on filesystems occupying partitions, but it also resizes filesystems that have been stored as single files by `ntfscdclone`.

Note that `ntfsresize` doesn't change partition tables, it changes only the NT filesystem inside the partition. Changing the partition table is a job for `fdisk` or `sfdisk`.

This article does not cover how to partition a disk. A detailed and cautious description of how to free space on a drive occupied entirely by a single NT filesystem could take an article at least as long as this one. The operation itself doesn't take long, but it is a bit dangerous. Carelessness, or even bad luck, could result in a computer that refuses to boot. Given this, and given that the workaround of an extra hard drive costs almost the same as a tank of gas, this article continues to assume that partitioning already has been done.

Suppose, however, that the NT partition is just a little too small for the NT filesystem. This can happen, for example, if you don't account for the need of most partitioning tools to round down to a nearby sector, or if you replaced a defective drive with one having the same advertised size but with a different geometry.

The `ntfscdclone` utility will work just fine on a partition that is too big, but it refuses to fit into a space that is even the slightest bit too small.

In that case, the `ntfsresize` tool can come to the rescue. To figure out how much space you could shrink out of your NT filesystem, type the command that follows (from the directory containing `ntfs.img`):

```
ntfsresize --info ntfs.img
```

The software will report something like the following:

```
ntfsresize v1.11.2
Device name       : ntfs.img
NTFS volume version: 3.1
Cluster size     : 4096 bytes
Current volume size: 90009203200 bytes (90010 MB)
Current device size: 90009203200 bytes (90010 MB)
Checking filesystem consistency ...
100.00 percent completed
Accounting clusters ...
Space in use     : 6508 MB (7.2%)
Collecting resizing constraints ...
You might resize at 6507421696 bytes or 6508 MB
(freeing 83502 MB). Please make a test run using both
the -n and -s options before real resizing!
```

This reports that you could shrink your filesystem down to as little as 6,508MB. Windows probably wouldn't run if you reduced it to the minimum size; it would be smart to leave a little room for future growth anyway. Note that the software advises that you could make a "test run using both the `-n` and `-s` options". Instead, you simply could keep a backup copy in a safe place in case something goes wrong. Or, you could do both. Shrinking the filesystem to 10,000MB requires the following command:

```
ntfsresize --size=10000M ntfs.img
```

This produces a great deal of feedback, including the following:

```
100.00 percent completed
Updating $BadClust file ...
Updating $Bitmap file ...
Updating Boot record ...
Syncing device ...
Successfully resized NTFS on device 'ntfs.img'.
```

This should create an NT filesystem small enough to fit into its designated partition.

Conclusion

The NTFS tools may not be a requirement for everyone wanting a secure Windows workstation, but they do make life a lot easier.

In the context of a single dual-boot computer, complete backups can be performed to a safe, non-NTFS partition, either on the same hard drive, or even onto a removable hard drive of sufficient capacity. This may not make the effort worthwhile for everyone. However, for the user already equipped with a dual-boot system, the tools for greatly enhanced security may already be installed.

For a network administrator in charge of many Windows workstations, the potential is even greater. Dual-boot computers can be equipped with a shared disk partition (see Kevin Farnham's article "The Ultimate Linux/Windows System" in the June 2006 issue of *Linux Journal*). If GRUB is installed in this shared partition, along with alternate menu files, scripts can be written that reboot the computer into runlevels that automatically restore the Windows image, update it and so on.

Windows and Linux may be competitors in many areas. However, one of the great strengths of Linux is its open nature and the versatility of its command-line tools. The Linux-NTFS tools open up a conversation with the NT filesystem that, because of its one-way nature, makes for ideal security. ■

Steven Mathes installed Linux on his computer for the first time in 1995, when it was possible to back up Windows with tar. He can be reached at smathes@tiac.net.

Break the Hardware Upgrade Cycle with

Win4Lin

WINDOWS VIRTUAL DESKTOP SERVER



Smooth Migration to Linux by delivering Windows
to Linux clients via Win4Lin VDS.

Jon Watson

Win4Lin Virtual Desktop Server (www.win4lin.com) is a client/server virtualization solution that can be used to migrate an organization from an expensive and high-maintenance Windows infrastructure to a more robust and sleeker Linux base gently. How? A single copy of Windows can be delivered directly to multiple users' desktops with the click of a mouse. You don't want your users to have a full Windows desktop? No problem. Virtual Desktop Server (VDS) can be configured to deliver a single application to the Linux desktop instead. Given its flexibility, it's not surprising that the reasons for using VDS to move to a Linux infrastructure are equally as varied.

Linux server requirements are generally lower than those of Windows servers. Therefore, Win4Lin VDS can be used to break the hardware upgrade cycle. With Vista on the horizon, many organizations are faced with potentially costly hardware upgrades in the next few years.

Although arguments can be made on either side of the Total Cost of Ownership (TCO) issue, organizations that have come to the conclusion that Linux offers a lower TCO are then faced with the technical and logistical burdens of migrating their infrastructure. VDS allows the baseline software swap to occur while still allowing employees to continue using their familiar Windows environment and applications.

VDS offers organizations indefinite breathing room. Once the OS baseline has been swapped out, organizations can choose to remain in the Linux/Windows VM posture, or they can carry on with the business of sourcing or porting Linux solutions to their functional applications.

VDS offers "single application" deployment to the desktop, meaning that single mission-critical apps need not ever be ported. Single Windows applications can be launched right into the client Linux desktop.

Aside from the infrastructure questions, running a VDS server allows for a centralized point of management, upgrades and maintenance. Because all clients are being served the same Windows image, single changes on the server end mean rapid organization-wide change.

Existing Windows licenses can continue to be used under their respective terms.

VDS is different from a lot of client/server virtual machine (VM) solutions on the market. Most VM server products simply provide a remote display to the client, whereas VDS provides a proper client/server X Window System display to the client. If the client cannot support X Protocol messages, the alternative "display" methods can be used via traditional Virtual Network Connection or Tarantella, which really opens up the door for pretty much any client with a recent Web browser installed on it.

One of the main tasks for many Windows administrators is keeping Windows patched and updated in order to protect clients from the many spyware and malware attacks perpetrated against hapless Windows machines on a daily basis. As mentioned, VDS offers a single instance of Windows to patch and upgrade, which not only takes less time, but also offers more simplicity than staging patches throughout the organization. Further, because the end users' environment is a product of a combination of the master Windows image and their own locally stored settings, simply logging off and logging back in refreshes their session with the master image and thus eliminates any running malware or spyware in their session.

Clients

Win4Lin recommends using the native Win4Lin Terminal Services Client in order to make use of all the advanced functionality a native Win4Lin client/server connection offers. However, there are a plethora of ways to connect to a VDS server, and unless you're desperately in need of seamless printing, almost any client will get the job done.

Connection to a VDS server is possible with Telnet, rlogin or SSH (with X11 forwarding enabled). For Telnet and rlogin, the `{DISPLAY}` environment variable must be set correctly. In general, the remote login options are suitable only for high-speed environments, such as local LANs. WAN and consumer-grade high-speed Internet connections do not generally provide enough bandwidth to use these methods. Connections also are possible using the RealVNC client, the NoMachine client and Tarantella.

In short, if you can't find a way to connect to the VDS, you're simply not trying.

Licensing

A base license starts at \$2,500 US for 25 seats. Bump licenses can be obtained in various increments to enable VDS to handle up to 1,000 users. Whether your server can handle 1,000 users is up to you to decide. It's important to note that these are licenses for VDS and not for Windows. Organizations will have to provide their own Windows licenses under Microsoft's conditions. In most cases, however, existing licenses can be used.

Download and Installation

The current version of Win4Lin VDS is 3.0, but 3.5 is in tail-end beta and probably will be released before this article goes to print. Because it doesn't make a whole lot of sense to write an article about a potentially moving target, we use the 3.0 stable version.

Win4Lin Pro and VDS are the same binary, but different licenses unlock different functionality. There are DEB and RPM packages for 32- and 64-bit Linux. The installation prompts for a license code, and that's when the product turns into either Pro or VDS.

There is no upgrade path from Win4Lin Pro to Win4Lin VDS, so if you have Pro installed, you'll have to relicense it with a VDS license by using the `ask_license.sh` script.

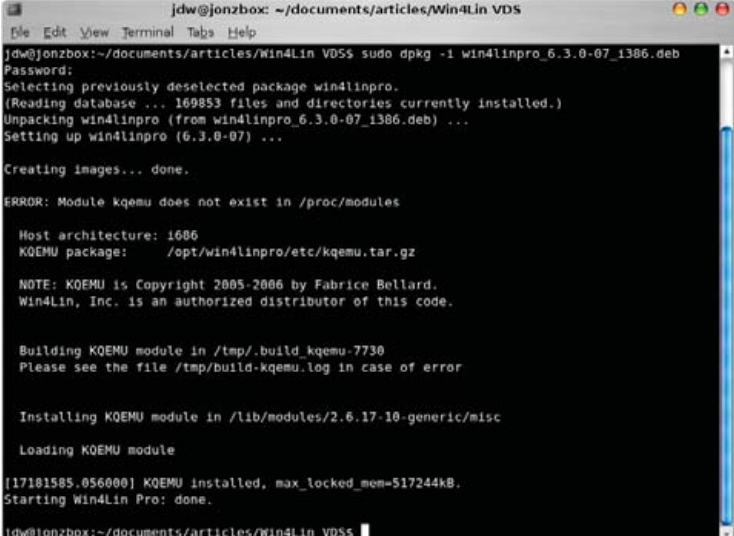
VDS clients are available for Linux, Solaris and Windows, but the source also is available for download, which allows moderately skilled Linux users to compile a client for almost any platform.

Install the Server

We downloaded our packages from the Win4Lin FTP site (ftp.win4lin.com/pub/releases/linux/pro/3.0/index.html). There are server packages available as 32- and 64-bit RPM, 32- and 64-bit DEB and tarball.

We downloaded the 32-bit RPM package for our test platform, a Centrino Duo Core running at 1.83GHz with 1GB of RAM running a beta of Edgy Eft. We ended up moving to SUSE 10.1 after the initial install, because Edgy's "edgy" kernel had issues with the Win4Lin client. Serves us right for trying to use a beta release on our testing platform, and we've been properly chastened. I mention it only because some sharp-eyed readers may be able to detect differing platforms in some of the screenshots.

The `win4linpro_6.3.0-07_i386.deb` package was only 3.7MB and installed without a hitch. The VDS installation manual advises that the toolchain for supported distros must be installed prior to installing VDS. The reasoning behind this is that all Win4Lin products provide a specialized KQEMU module in order to deliver satisfactory speed performance. As long as you've dutifully installed the toolchain, the building and insertion of said module will be largely transparent to you (Figure 1).



```
jd@jonzbox: ~/documents/articles/Win4Lin VDS
File Edit View Terminal Tabs Help
jd@jonzbox:~/documents/articles/Win4Lin VDS$ sudo dpkg -i win4linpro_6.3.0-07_i386.deb
Password:
Selecting previously deselected package win4linpro.
(Reading database ... 169853 files and directories currently installed.)
Unpacking win4linpro (from win4linpro_6.3.0-07_i386.deb) ...
Setting up win4linpro (6.3.0-07) ...

Creating images... done.

ERROR: Module kqemu does not exist in /proc/modules

Host architecture: i686
KQEMU package: /opt/win4linpro/etc/kqemu.tar.gz

NOTE: KQEMU is Copyright 2005-2006 by Fabrice Bellard.
Win4Lin, Inc. is an authorized distributor of this code.

Building KQEMU module in /tmp/build.kqemu-7730
Please see the file /tmp/build-kqemu.log in case of error

Installing KQEMU module in /lib/modules/2.6.17-10-generic/misc

Loading KQEMU module

[17181585.056000] KQEMU installed, max_locked_mem=517244kB.
Starting Win4Lin Pro: done.
jd@jonzbox:~/documents/articles/Win4Lin VDS$
```

Figure 1. Building and Installing the Modules

Now it's time to install my single Windows instance. We used Windows XP Home, installed it and allowed it to upgrade itself to service pack 2 before provisioning accounts. To do so, we put our Windows XP CD into the drive and ran the command:

```
sudo loadwinproCD
```

If you haven't installed your VDS license yet, you will be prompted to do so at this time. You can choose not to enter a license at this time and use the product for 14 days, but none of the server functionality will be enabled, effectively leaving you with a single-user workstation.

In Figure 1, you will see that we had to use the `-r` switch to load the guest media. This is because our testing bench already had Win4Lin Pro installed on it previously, and the `loadwinproCD` command detected this. The `-r` switch simply tells VDS to reload the guest media.

There are two steps to installing VDS. First, the superuser must copy the guest media to the hard disk, which is what `loadwinproCD` does. The next step is actually to install Windows, which is done under a normally privileged user account. The process of installing Windows and creating the master profile, from which all other user accounts will derive their Windows environment, are tightly coupled—so much so, in fact, the entire next section is devoted to understanding the process.

Profile-Based Provisioning

Win4Lin VDS uses profile-based provisioning. In a nutshell, this means that a single master profile must be configured. This is the profile all other users will inherit from, and this is also the only profile where patches, upgrades and applications need be installed and other maintenances need be performed. Once a master profile has been configured as desired, individual user profiles are created for each system user.

Create a Master Profile

VDS must be installed under a user account. As with all infrastructure decisions, a few moments spent now can save countless headaches later. I've decided to install my master profile under my main user account `jdw`. I've further decided that I'm going to provision two other user accounts on my system named `jwtatson` and `dwatson`. My first step, then, is to create the two `jwtatson` and `dwatson` accounts. Organizations already running a Linux

system quite likely have all of their user accounts set up and may be required to create only an account for the master VDS profile.

To create the master profile, first I log in to my system and run:

```
sudo adduser jwtatson
sudo adduser dwatson
```

After assigning these new user account passwords, they are ready to go. Now it's time to create the master profile. This involves installing Windows, designating that installation as the master and then configuring it as required.

So, I log out and back in as my master profile user `jdw`, and then run the `installwinpro` command.

Various options to this command dictate the characteristics of the virtual machine into which Windows will be installed. I accept the defaults, which include a 4GB disk image size. Depending on the amount of applications you intend to install, 4GB may not be enough. There is no need to take user space into account, however, because individual user's documents and settings are stored in the Linux filesystem on a user's respective workstation and not within the image (Figure 2).

The default profile name will be `winpro` unless you change it with the `-d` switch. Keep in mind that if you switch the configuration name, you will need that information to export the master profile in the next step.

Take my word for this—back up your image now. Living through one Windows install is painful enough, you don't want to have to do it twice.

Backing up your image is as simple as copying the `GUEST.IMG` file to another location. If you took the default installation options, you'll find the `GUEST.IMG` file in `/home/jdw/winpro/` (obviously, you'll need to substitute your own user name). If you modified your installation, you'll know where to find it.

Once installation is complete, designate this installation as the master by running the command `/opt/win4linpro/bin/export-profile <configuration name>`.

If you changed the configuration name while installing the guest session in the previous step, you need to provide that name in place of `<configuration name>`. If you didn't specify a configuration name, the default `winpro` was used, and you need not supply it to export the master profile.

Launch the image and customize it. It is critical that you launch and shut down the image at least once in order to create the master profile properly. To launch the master Windows profile, either double-click the Win4Lin icon on your desktop, or run the command `winpro` from the command line.

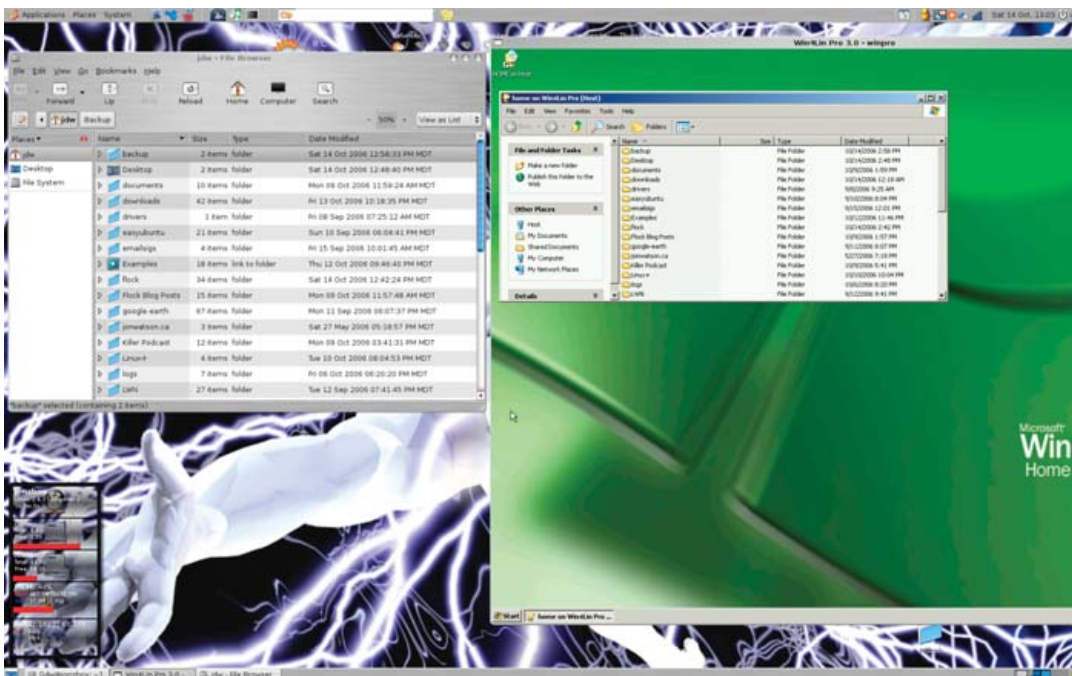


Figure 2. User documents and folders are stored on the Linux filesystem.

Customizing the Image

Because provisioned users will receive a fresh copy of the master image each time they log in to the VDS, it's not critical that you fully customize your image right away. In theory, you can skip to creating your user profiles right now and then come back and customize your master image later. In practice, however, there are a few reasons why you probably should do it now:

- Familiarity: it's quite likely that one of the reasons your organization is using VDS to swap out the infrastructure is to minimize the impact on users. Therefore, common sense seems to say that every effort should be made to provide users with the most familiar desktop and set of applications possible.



Figure 3. Update your Windows master image, but lock down a known good state.

- Security: because users' sessions are refreshed with a copy of the master image every time they log in to the VDS, it's easy to think that security can take a back seat. If the OS is refreshed back to the master image at least daily, how much damage can spyware, malware or a virus do? That's a good question, but consider what can happen if your master image contains malware or a virus. Lock it down, now (Figure 3).
- Technical: the master profile cannot be running when any user profiles are running. Because you have to fire up the master profile in order to make changes to it, failing to customize it now might mean a lot of off-hours work in the near future—or significant work disruptions as all users are forced to log off and stay off while the master image is worked on.

Now is probably a good time to take another backup of your image. Every time you make a change to the master profile, it's wise to create another backup of the image. Remember that if your image becomes corrupted, a whole lot of users won't be able to get their work done until you've restored it. It takes only a few minutes to copy a 4GB image back to the master directory, but it takes a lot longer to re-install Windows or bring an old image up to date.

The Win4Lin VDS service should start itself. If it doesn't, or if you need to restart it from some reason, you can do so with:

```
/opt/win4linpro/etc/mergepro_rc start
```

or:

```
/opt/win4linpro/etc/mergepro_rc restart
```

Create User Profiles

Only provisioned users will be able to use the master profile. System users that are not intended to use VDS need not be provisioned, but the provisioning process must be done under each user's account. If you have only a few users, it might be quicker to `su` to each user account and run the import command manually. If you have many users, however, some clever bash scripting or log-on scripts might be in order to facilitate the process. In my example, this means I have to log in as both `jwtatson` and `dwatson` and run the following command:

```
/opt/win4linpro/bin/import-profile /home/jdw/winpro
```

The master profile has been created, and users have been provisioned. The bulk of the server work is done now, and any changes made to the master profile from this point on (such as new applications being installed or patches being applied) will propagate on down to provisioned users each time they log on.

But, how will these users log on? It's time to install a client.

Install the Client

As mentioned before, there are several ways to connect to a VDS. I look only at the Win4Lin client as it is free and easily attainable.

Using the native Win4Lin client against the VDS server provides the best speed and feature set. The Win4Lin client can be downloaded from the Win4Lin site (www.win4lin.com/component/option,com_repository/Itemid,76/func,fileinfo/id,2) and comes in flavours for Linux, Solaris, Windows and source.

Strangely, although the VDS itself is available in both DEB, RPM and a

DEDICATED SERVERS

ROCK BOTTOM PRICING



\$60^{mo}

Celeron D

\$160^{mo}

Dual Core Xeon

- » 1300GB/mo Bandwidth
- » 1GB RAM DDR400!
- » 160 GB SATA2 Hard Drive
- » 100Mbps Dedicated Port

- » FREE REBOOTS 24x7
- » FREE SUPPORT 24x7
- » FREE SETUP Available*
- » BRAND NAME Hardware

w.cari.net

* See website for details.

carinet
betterServers. betterService.




Figure 4. The Win4Lin Client Running a Desktop

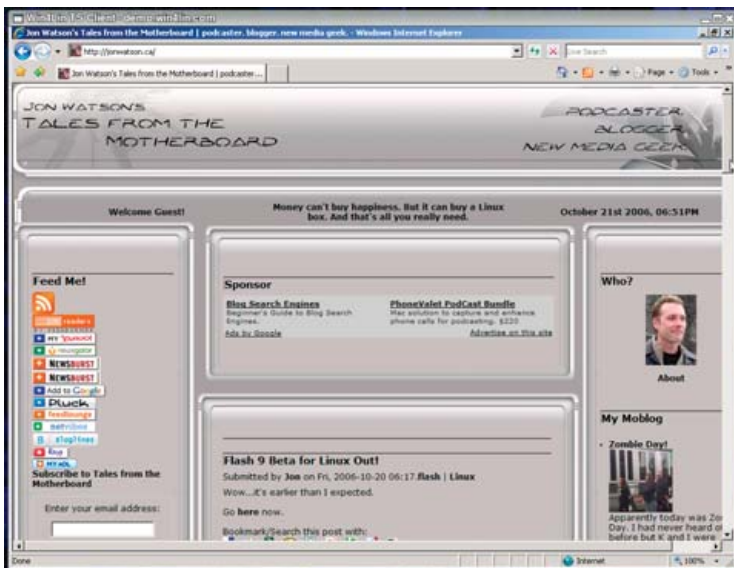


Figure 5. Win4Lin can deliver only a single application instead of the entire desktop.

tarball, the Linux client isn't available in DEB format. Because installing from source usually makes me lose my lunch, we're going to grab the RPM, use alien to convert it to a DEB and then use dpkg to install it on a Debian-based system. Here are the steps:

1. Download the RPM from Win4Lin.
2. Run `sudo alien wtsclient_1.0.0-4_i386.rpm`.
3. Run `sudo dpkg -i wtsclient_1.0.0-4_i386.deb`.
4. Run the `wtsclient` command to connect to the server.

The Win4Lin VDS can be configured to deliver a single application or an entire Windows desktop. We connected both configurations to the Win4Lin demo server and provided two screenshots. The first

Because all clients are being served the same Windows image, single changes on the server end mean rapid organization-wide change.

screenshot shows an entire Windows desktop configuration, and the second shows only Internet Explorer being delivered to our Linux desktop (Figures 4 and 5).

Delivering Single Applications to the Linux Desktop

So far, all we've covered is how to get a full Windows XP desktop delivered to a Linux desktop. Although that might be suitable for many situations, there are others when users may require only a single Windows application. How can VDS be rigged to open up an application instead of a full desktop? By tweaking the Windows registry, of course. The Win4Lin VDS manual is the best place to look for current instructions on how to achieve this, but in the interests of sitting down and getting it done with just this copy of *Linux Journal*, now we provide steps required to deliver a single application.

Ensure Correct Win4Lin Registry Key Value (All Versions of Windows)

Regardless of the application or the version of Windows being used, a Win4Lin registry key must be set or verified first:

1. Open the regedit application.
2. Navigate to the `HKLM\Software\Microsoft\Windows\NTCurrentVersion\Winlogon`.
3. Ensure the `Userinit` variable reads exactly `B:\mrgpro32.exe`. If the value isn't exact, change it.

It doesn't make a whole lot of sense to have users log in to Windows just to run a single application. Therefore, step one—although optional—is to set the master Windows profile to log in a user automatically. Different flavours of Windows have different provisions for allowing automatic log in.

Set Autologin (Windows 2000)

1. Launch the Control Panel.
2. Launch the Users and Passwords applet.
3. Uncheck the box that reads Users must enter a user name and password to use this computer, and click OK.
4. When prompted, enter the user name and password of the account under which you would like Windows to launch.

Set Autologin (Windows XP)

1. Launch the Control Panel.
2. Launch the User Accounts category.
3. Click Change the way users log on/off.
4. Uncheck the Use the Welcome screen check box.

- Click Apply Options.
- Launch the alternate user account editor by clicking Start→Run and entering control userpasswords2. Click OK.
- Uncheck the box that reads Users must enter a user name and password to use this computer. Click OK.
- When prompted, enter the user name and password of the account under which you would like Windows to launch.

Designate the Single Application to Deliver (All Versions of Windows)

- Launch the registry editor by clicking Start→Run. Type in regedit and click OK.
- Navigate to HKLM\Software\Win4Lin.
- Right-click in the empty right-hand pane to create a new variable.

- Select String Value.
- Type SingleAppStart (case-sensitive).
- Double-click on the newly created SingleAppStart variable.
- In the Value data: field, enter the full path to the executable to be launched. For example, to run Microsoft Word, WORD.EXE is not sufficient. The full path of C:\Program Files\Microsoft Office\word.exe (or wherever your Word executable is located and named) is required.
- Exit the registry editor, and you're done!

Now when users launch their clients and log in, Microsoft Word will launch onto their desktop right beside their Linux applications. This may not seem like a big deal, because the Microsoft Office suite is nicely supported by Wine and CrossOver Office, but swap out Word for absolutely any other application on your Windows desktop, and the power becomes obvious. Because a full copy of

Windows is being brought to bear to deliver the application, there is an almost unlimited number of Windows applications that can be delivered in this manner with zero modification.

Parting Thoughts

Virtualization technology isn't new. IBM has been playing with it since the 1960s. What's making virtualization exciting again is the widespread availability to fast networks and powerful servers. Using Linux to deliver Windows is cost effective in terms of hardware, management and training, and products such as Win4Lin Virtual Desktop Server make the technology easy to install and use. In fact, virtualization technology has come so far that the tricky points are no longer technical in nature; they are logistical. It's more difficult to plan a virtualization strategy than it is to implement one. ■

Jon Watson (www.jonwatson.ca) is a Canadian GNU/Linux enthusiast who regularly contributes articles to the Linux community. When not writing, blogging and podcasting about free and open-source software, Jon frequently can be found in his office polishing his Linux+ certification, which impresses no one but himself.

SilverComet SZ

- Sony VAIO SZ
- 13.3" WXGA+ w/ X@1280x800
 - NVidia GeForce Go graphics
 - 1667–2167 MHz Core 2 Duo
 - 512 MB–2 GB RAM
 - 80–160 GB hard drive
 - CDRW/DVD or DVD±RW:DL
 - 3.7–4.1 pounds
 - 10/100 Mbps ethernet
 - 802.11a/b/g (54Mbps) WiFi
 - Starts at \$1850




Rhino D820/M90

- Dell Latitude D820/Precision M90
- Up to 17" WUXGA w/ X@1920x1200
 - NVidia Quadro FX 3500M graphics
 - 1667–2333 MHz Core 2 Duo
 - 512 MB–4 GB RAM
 - 40–160 GB Hard Drive
 - CDRW/DVD or DVD±RW
 - 6.3–8.6 pounds
 - 10/100/1000 Mbps ethernet
 - 802.11a/b/g (54Mbps) WiFi
 - ExpressCard/EVDO
 - Starts at \$1455



Toucan T60/T60ws

- ThinkPad T60/T60ws by Lenovo
- Up to 15.4" WSXGA+ 1680x1050
 - ATI Mobility FireGL V5200
 - 1833–2333 MHz Core 2 Duo
 - 512 MB–4 GB RAM
 - 60–120 GB hard drive
 - CDRW/DVD or DVD±RW
 - 5.2–6 pounds
 - 10/100/1000 Mbps ethernet
 - 802.11a/b/g (54Mbps) WiFi
 - Bluetooth
 - Starts at \$1950



EmperorLinux

...where Linux & laptops converge

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

Model specifications and availability may vary.



What do we do?
Get the details
on page 15
of this issue



Accessing Linux Filesystems in Windows

You don't have to wait for Microsoft to support Linux filesystems to get at your Linux files from Windows.

IRFAN HABIB

Illustration by Luis Carlos Torres

Lack of access to your data in a new operating system may be one of the most severe impediments for doing an OS migration. There is little personal incentive for users to switch to a system that can't interoperate with their data, as the system would be practically useless to them.

Linux has done a great job in allowing Microsoft Windows users to access their Windows partitions from Linux. Support for a majority of Windows filesystems is available, such as seamless support for FAT16/32 and partial but increasingly complete support for NTFS. There also are some tools, such as Captive-NTFS, which enable complete support for NTFS drives from Linux. Data access is not restricted to the local host.

Samba allows Linux users to access their shared data on Windows computers over a network. Hence, we can say that, for MS Windows users, access to their data is no longer an impediment to Linux migration.

However, as Linux advances into the desktop, many people use Linux for some tasks and then turn to Windows for software that is not yet available on Linux, such as high-end games, Adobe Photoshop and various domain-specific applications for which no open-source equivalent exists. When people are using such applications, they generally require access to their Linux partitions. Support for Linux filesystems is non-existent in Windows. Thus, Linux dual-booters must use some tools to access their Linux filesystems.

Tools for Accessing Linux Partitions in Windows

As mentioned previously, Windows does not have native support for Linux filesystems. All is not lost, however. The Open Source community has risen to the challenge and created some excellent software to solve this problem. This article focuses mainly on LTOOLS, which is advanced software with multiple interfaces that allows users to access a range of Linux filesystems. But first, let's skim through some other existing software that could do the task.

Ext2fsd

Ext2fsd is one of the oldest projects in this area. It allows access from Windows to ext2 filesystems and can be downloaded from sourceforge.net/projects/ext2fsd. It installs as a filesystem driver, not as a regular application. Making Ext2fsd a filesystem driver integrates ext2 partitions transparently into Windows and allows Windows to use ext2 partitions as if they were a native format and enables full read/write support. Ext2fsd does not only read ext2 partitions, but also Ext2 was one of the first de facto Linux filesystems, and many new Linux filesystems, such as ext3 and ext4 are backward-compatible with it. Thus, the driver can work with ext3 and possibly ext4. The picture is currently unclear with respect to ext4, as ext4 was recently added to the mainline kernel for testing. But, when using ext3/4 with Ext2fsd, you will be using only the features of ext2; any other additional features, such as enhanced journaling capabilities, will not be used.

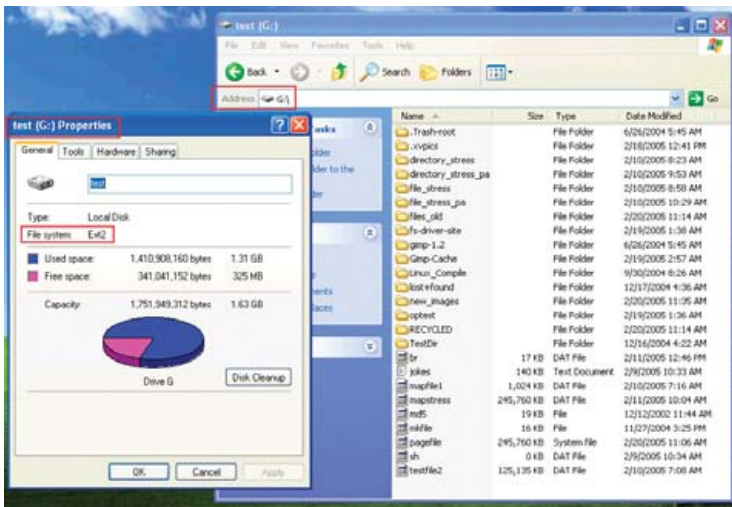


Figure 1. Ext2fsd makes an ext2 filesystem look like any other filesystem in Windows.

rfstool

ReiserFS increasingly has become a popular Linux filesystem, because of its fault-tolerance capabilities. rfstool allows access to ReiserFS partitions from Windows to Linux; however, it supports only read-only access, and the developers, according to their Web site, have no plans to change that. The tool is available from freshmeat.net/projects/rfstool.

LTOOLS

The previous tools lead us to the tool to which this article is dedicated. Unlike Ext2fsd and rfstool, which are specific to one particular class of filesystems, LTOOLS are more generic. They support ext2, ext3 and ReiserFS. LTOOLS are a set of command-line tools, along with two GUIs and a Web-based front end, to enable the reading of and writing to Linux ext2/3 and ReiserFS filesystems from nearly all DOS or Windows (XP, 2000, NT, ME, 9.x or 3.x) versions, running on the same machine or remotely. So, whenever you're running DOS or Windows, and you desperately need to read or write to a Linux partition, which may be on your own computer or any other, you can make use of LTOOLS. LTOOLS also is a great tool for fixing your Linux installation, if you do not have a live CD.

A SHORT PRIMER ON FILESYSTEMS

Some readers might be wondering what a filesystem really is. A filesystem basically defines a method for storing and retrieving files from a disk. This begs the questions: "Why are there so many filesystems? Why can't everyone decide on the best way to store and retrieve data from a disk and make that filesystem standard across all platforms?" Different operating systems come with different filesystems, because they target different users. Windows NT, which was targeted at enterprise users, came with a filesystem called NTFS, which was a filesystem that had enhanced security; whereas Windows 9x, which was targeted at ordinary users, came with FAT16/32, which was a filesystem with less security but with more performance. Windows XP offers both, as it is marketed to both segments. Similarly, on Linux, ext2/3 are the de facto filesystems, but there are some others for more special-purpose uses, for example, filesystems for high-performance computing, such as XFS, or filesystems with a great deal of fault tolerance, such as ext3 and ReiserFS.

Adding "support" to an OS for a particular filesystem basically entails that we define the data structures involved to the operating system upon which the data is stored on the disk. It is more difficult to add support for proprietary filesystems, such as NTFS, because the structure of the data on the disk, the encryption algorithms and so forth, are not known to the public. That is why it is proving to be a challenge to support NTFS completely in Linux.

As mentioned previously, LTOOLS comes with two different GUI interfaces to enable you to access your Linux partitions. LTOOLS comes with LTOOLSgui, which is a Java-based graphical user interface for local or remote access to your Linux files, and LTOOLSnet, which is a Microsoft .NET-based user interface, which also provides local or remote access.

If you do not like using non-free Java or MS .NET, you can use your Web browser as a graphical front end for LTOOLS. To achieve this functionality, the package contains LREADsrv, which is a simple Web server, making your Linux filesystem available in an Explorer-like view in your Web browser. Using LREADsrv, you can allow remote access to your Linux partitions, as well as to your DOS/Windows partitions.

Installation of LTOOLS

LTOOLS comes with a default Windows installer, which seems quite dated. After following the normal installation procedure, the installer creates an entry in your start menu called LTOOLS from which you can access a plethora of interfaces that allows you to access your Linux filesystems.

LTOOLS comes with support for nearly all Windows versions; however, all of the interfaces won't run on all Windows versions. LTOOLS provides two different console versions for Win9x/ME and Windows NT/XP.

Overview of LTOOLS—Command-Line Interface

The command-line interface provides basic functionality for writing and retrieving data from Linux. LTOOLS commands have the following format. All commands have three files associated with them, for example:

```
ldir.bat
```

This command lists directories; however, it is not a program but a script. This script, depending on your system, further invokes either of these two programs: LdirDOS.exe or LdirNT.exe. The first one is for Win9x/ME, and the second one is for Windows NT/XP.

Many LTOOLS commands have a logical syntax. For example, partition names are Linux names. So, if I want to copy a file called vars.inc from /root (which is on /dev/hda2) to my C: drive, I would do the following:

▶ FEATURE: ACCESSING LINUX FILESYSTEMS IN WINDOWS

```
lread.bat --s/dev/hda2 /root/vars.inc C:\vars.inc
```

Similarly, for writing to Linux, I would do this:

```
lwrite.bat --s/dev/hda2 C:\vars.inc /root/vars.inc
```

Along the same lines, LTOOLS also has the commands shown in Table 1.

ReiserFS is not supported via the above-mentioned tools. Thus, LTOOLS also ships with `rfstool`, which can be used to read from ReiserFS partitions. In order to read the hard disk under Windows NT/2000/XP or UNIX/Linux, you need administrator rights. If you are running LTOOLS under a non-administrator account, you may not be able to access the hard disk. LTOOLS does not respect Linux ownership. This means that if users were to mount the root device, they could change anything, including `/etc/passwd/`.

LTOOLS .NET Framework Client

The .NET framework client is one of the most feature-rich clients available in LTOOLS. To run it, you need to download a copy of the Microsoft .NET framework from the Microsoft Download Web site (www.microsoft.com/downloads).

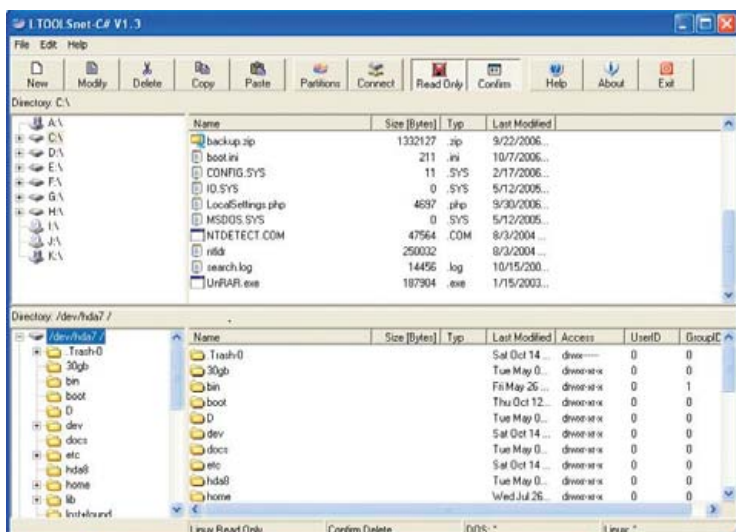
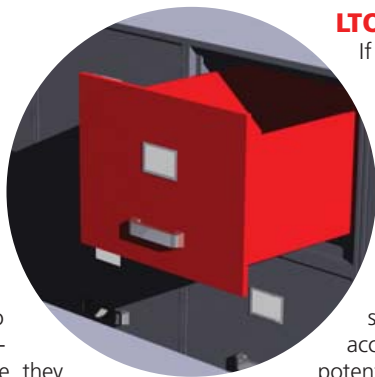


Figure 2. The C# LTOOLS File Manager

The client allows you to view all Windows and Linux partitions, and you can transfer files between them, delete files, edit files and modify them. It is also possible to mount a remote device and edit its contents. This is extremely useful when I have some problems with my Web server. I mount the drive remotely if I'm using a Windows machine and make all

the modifications necessary to get it running.



LTOOLS Java Framework Client

If your Windows installation is Windows95/98/ME and does not support the MS .NET framework, the Java interface is for you. To run the Java interface, you need a copy of the Java Runtime Environment, which you can download from java.sun.com. The Java interface has features analogous to the .NET client.

LTOOLS Web-Based Interface

The best interface in LTOOLS, based on my experience, is the Web-based service. LTOOLS comes with a built-in Web server, `LREADsrv.exe`, which allows users to start it and access their filesystems via a Web browser. This has great potential if you want to share files with other people remotely. I would not recommend running `LREADsrv.exe` on a server that is globally accessible, as it could compromise your data, so you should share it in an environment where only legitimate users have access to it, such as a virtual private network. `LREADsrv.exe` still has some problems; however, they will be fixed in future releases.

`LREADsrv` is still alpha and has certain limitations, which include problems with HTTP 1.1 Web browsers, such as Internet Explorer, which slows the response from the server considerably. Another limitation is that `LREADsrv`, in its current version, has been implemented as a mono-threaded application—meaning that if multiple people are accessing the filesystem at the same time, the changes they make are applied globally, which can lead to lost updates and concurrency problems. `LREADsrv`'s error checking is weak. Most user input (filenames and so on) is not validated. So, if a user types some filenames incorrectly or mistypes a hard disk partition, the Web server can go into an unstable state, which, fortunately, does not result in any data loss.

Conclusion

Linux users have increasingly more mature support for Windows filesystems. LTOOLS provides a unified way to access the most popular Linux filesystems through a plethora of interfaces from Windows. However, support for Linux filesystems in Windows still has a way to go. Windows support for various other Linux/open-source filesystems, such as XFS, is not yet available. Drivers capable of using advanced features, such as journaling in `ext3` and `ReiserFS`, are not mature. Integration of Linux filesystems with Windows is an important area, and the lack of it can be a serious impediment to an OS migration. Thus, to enable enhanced interoperability between MS Windows and Linux, given that Windows is still the dominant desktop operating system, the Open Source community must focus on adding mutual support for filesystems. ■

Irfan Habib is an undergraduate student of software engineering at the National University of Sciences and Technology, Pakistan. He has been deeply interested in Free and Open Source software for years. He often works in heterogeneous computing environments—that's why mutual support for filesystems of different platforms is important to him.

Table 1. LTOOLS Commands

<code>lread.bat</code>	Read and copy files from Linux to DOS; sample usage: <code>lread.bat --s/dev/hda2/root/vars.inc C:\vars.inc</code>
<code>lwrite.bat</code>	Copy files from DOS to Linux; sample usage: <code>lwrite.bat --s/dev/hda2 C:\vars.inc/root/vars.inc</code>
<code>ldel.bat</code>	Delete Linux files or (empty) directories—same as <code>rm -f</code> and <code>rmdir</code> in Linux; sample usage: <code>ldel.bat /root/vars.inc</code>
<code>lchange.bat</code>	Change Linux file attributes and owners—analogueous to <code>chmod</code> ; sample usage: <code>lchange.bat --s/dev/hda2 754 /root/vars.inc</code>
<code>lren.bat</code>	Rename Linux files—analogueous to <code>mv</code> ; sample usage: <code>lren.bat --s/dev/hda2/root/vars.inc/root/var2.inc</code>
<code>lmkdir.bat</code>	Create a new Linux directory—analogueous to <code>mkdir</code> ; sample usage: <code>lmkdir.bat --s/dev/hda2/root/newdir</code>
<code>lln.bat</code>	Create a symbolic link—analogueous to <code>ln</code> ; sample usage: <code>lln.bat --s/dev/hda8 /root/link/root/vars.inc</code>
<code>lcd.bat</code>	Change directory—analogueous to <code>cd</code> ; sample usage: <code>lcd.bat /home/</code>
<code>ldrive.bat</code>	Set the default Linux disk drive; sample usage: <code>ldrive /dev/hda8</code>

Polywell Ultimate Linux Systems

More Choices, Better Service, Best Value



Tower Systems

4-way Opteron, Quadro Graphics PolyStation 2055A-2210

- 2xAMD Opteron™ Dual-Core 2210 Processors
- 2GB Dual Channel 667MHz DDR2 ECC Memory
- NVIDIA Quadro™ FX 1300 OpenGL Graphics
- 2x250GB SATA Drives, On-board RAID-5 Controller
- Sony 16X DVD-RW, Logitech K/B, Mouse
- 2x Gigabit LAN, 2 x PCIe 16x Slots
- Linux / BSD / Windows OS Supported

\$1999 (custom configuration available)

4-way Opteron, SAS Drives PolyStation 2055SS-2216

- 2xAMD Opteron™ Dual-Core 2216 Processors
- 8GB Dual Channel 667MHz DDR2 ECC Memory
- NVIDIA Quadro™ FX 3500 OpenGL Graphics
- 2x15K RPM 73GB SAS Drives, SAS Controller
- Sony 16X DVD-RW, Logitech K/B, Mouse
- 2x Gigabit LAN, 2 x PCIe 16x Slots
- Linux / BSD / Windows OS Supported

\$4695 (custom configuration available)

4-way Opteron, 32GB RAM Tower Server 2500M-2220

- 2xAMD Opteron™ Dual-Core 2220 Processors
- 32GB Dual Channel 667MHz DDR2 ECC Memory
- 11-Bay Tower with 2x700W Power Supply
- 3TB 6x500GB SATA Drives, RAID-5 Controller
- Sony 16X DVD-RW, Logitech K/B, Mouse
- 2x Gigabit LAN, 2 x PCIe 16x Slots
- Linux / BSD / Windows OS Supported

\$9500 (custom configuration available)

Desktop Systems

Low-Cost Linux Appliance PC Poly 485Ax-2800SP

- AMD Sempron™ 2800+ 64-bit Processor
- 256MB 667MHz DDR2 Memory
- ATI X1000 PCIe Graphics
- 80GB SATA Drive, SATA-RAID Controller
- 16X DVD-ROM Drive, Logitech K/B, Mouse
- 100Mbit LAN, 2 PCI + 1 PCIe 16x Slots
- Linux / BSD / Windows OS Supported

\$299 (for volume purchase only)

2TB Storage Appliance MiniBox 430AM2-3000SP

- AMD Athlon™ 64 3200+ 64-bit Processor
- 512MB Dual-Channel 667MHz DDR2 Memory
- NVIDIA GeForce™ 6150 Graphics
- 2TB 4x500GB SATA Drives, RAID-5 Controller
- Optional DVD or RW Drive, K/B, Mouse
- Gigabit LAN, 2 PCI + 1 PCIe 16x Slots
- Linux / BSD / Windows OS Supported

\$1295 (custom configuration available)

High-End PC, 4G DDR2 Poly 430AM2-FX62

- AMD Athlon™ 64 Dual-Core FX-62 Processor
- 4GB Dual Channel 667MHz DDR2 Memory
- NVIDIA GeForce™ 7900GS Graphics
- 10K RPM 74GB + 320GB SATA Drives
- 16X DVD-RW, 20-in-1 Media Reader, 1394a
- Gigabit LAN, 2 PCI + 1 PCIe 16x Slots
- Linux / BSD / Windows OS Supported

\$2250 (custom configuration available)



Rack Servers

Low-Cost ISP RackServer 1102EV-17 485Ax-3000SP

- AMD Sempron™ 3000+ 64-bit Processor
- 512MB 667MHz DDR2 Memory
- ATI X1000 PCIe Graphics
- 80GB SATA Drive, SATA-RAID Controller
- 1U 17" Short Rack, 1 Slim-CD, 2 x 3.5" Bays
- 100Mbit LAN, 1 Optional PCIe or PCI RISER
- Linux / BSD / Windows OS Supported

\$399 (for volume purchase only)

2GB RAM, 500GB ISP Server 1102EV-17 430AM2-4200

- AMD Athlon™ 64 X2 Dual-Core 4200+ Processor
- 2GB Dual-Channel 667MHz DDR2 Memory
- NVIDIA GeForce™ 6150 Graphics
- 500GB 2x250GB SATA Drives, RAID-5 Controller
- 1U 17" Short Rack, 1 Slim-CD, 2 x 3.5" Bays
- Gigabit LAN, 1 Optional PCIe or PCI RISER
- Linux / BSD / Windows OS Supported

\$999 (custom configuration available)

2GB ECC, 1TB ISP Server 1102EV-17 1000SL-1210

- AMD Opteron™ Dual-Core 1210 Processor
- 4GB Dual Channel 667MHz DDR2 ECC
- Integrated Graphics, 2 x Gigabit LAN
- 1TB 2x500GB SATA Drives, RAID-5 Controller
- 1U 17" Short Rack, 1 Slim-CD, 2 x 3.5" Bays
- Optional PCIe or PCI RISER Slot
- Linux / BSD / Windows OS Supported

\$1799 (custom configuration available)



High-Density Multi-Processor Servers



16-way Opteron, 128GB RAM 5124AIS 8850S-8212

- 8 x AMD Opteron™ Dual-Core 8212 Processors
- 128GB Dual Channel 667MHz DDR2 ECC Memory
- 12TB 24x500GB SATA Drives, RAID-5 Controller
- 4 x Gigabit Ethernet, 4 PCI-X, 2 PCIe Slots
- 5U 24-Bay RackCase, Redundant Power Supply
- Linux / BSD / Windows OS Supported

16-way 8212, 128G RAM, 24x500G **\$99999**
8-way 8212, 32GB RAM, 8x500GB **\$15999**



18TB 4U Storage, 32GB RAM 4024AIS 2500M-2210

- 2 x AMD Opteron™ Dual-Core 2210 Processors
- 32GB Dual Channel 667MHz DDR2 ECC Memory
- 18TB 24x750GB SATA Drives, RAID-5 Controller
- 2 x Gigabit Ethernet, 4 PCI-X, 2 PCIe Slots
- 4U 24-Bay RackCase, Redundant Power Supply
- Linux / BSD / Windows OS Supported

18TB 24x750G, 32G, 2x2210 **\$19999**
11TB 22x500G, 2GB, 1x2210 **\$9999**



6TB 2U Storage, 16GB RAM 2012SAS 2500M-2210

- 2 x AMD Opteron™ Dual-Core 2210 Processors
- 16GB Dual Channel 667MHz DDR2 ECC Memory
- 6TB 12x500GB SATA Drives, RAID-5 Controller
- 2 x Gigabit Ethernet, 4 PCI-X, 2 PCIe Slots
- 2U 12-Bay RackCase, 3 RISER or 7 Lowprofile slots
- Linux / BSD / Windows OS Supported

6TB 12x500G, 16G, 2x2210 **\$7999**
1.5TB 6x500G, 2GB, 1x2210 **\$3399**



2TB 1U Server, 32GB RAM 1104SC 2500M-2210

- 2 x AMD Opteron™ Dual-Core 2210 Processors
- 32GB Dual Channel 667MHz DDR2 ECC Memory
- 6TB 4x500GB SATA Drives, RAID-5 Controller
- 2 x Gigabit Ethernet, 1 PCI-X or PCIe RISER Slot
- 1U RackCase, 1 Slim CD, 4 Swap HD Bays
- Linux / BSD / Windows OS Supported

2TB 4x500G, 32G, 2x2210 **\$7888**
1TB 4x250G, 2GB, 2x2210 **\$1999**

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

Polywell OEM Services, Your Virtual Manufacturer

Prototype Development with Linux/FreeBSD Support
Small Scale to Mass Production Manufacturing
Fulfillment, Shipping and RMA Repairing



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

888.765.9686
www.Polywell.com/us/LJ



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

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

The Citadel Groupware Server



Microsoft Exchange, Meet Your Replacement

Jon Watson

The time has come to consider moving that expensive, high-maintenance Windows system to a sleeker, more robust Linux system. The gap analyses have been done, the meetings held, the presentations complete, and now it's go time. Although installing and configuring a Linux server back end can be challenging, we all know that users aren't going to care about that. What they want is uninterrupted functionality so they can continue doing their jobs. Although migrating users from applications such as Microsoft Office to OpenOffice.org is generally an intuitive task, the 800-pound gorilla that's keeping you up at night is e-mail and groupware. How are you going to provide and manage Microsoft Outlook-like functionality to the masses? In a word, Citadel.

One of the understated wonders of the Free and Open Source Software world is the Citadel Groupware Server (www.citadel.org). Controlled by a single developer, Citadel started life in 1987 as a UNIX version of the already-existing Citadel-CP/M application. Almost 20 years later, the modern-day Citadel boasts all of the functionality of a mature groupware server. One of the miracles that Citadel can perform is providing all of the most-used functionality of Microsoft Exchange with little fuss and even less cost.

Many modern organizations are coming to the realization that their IT budget is largely controlled by Microsoft's licensing fees and hardware requirements. Although organizations can prepare for some of these costs, many are looking at Vista with trepidation. Although the hardware requirements for Vista aren't obscenely over the top, many organizations still will need to upgrade their hardware in order to run it. And, sooner or later, run it they will. The hardware upgrade cycle is a never-ending source of pain for some organizations, because not only do servers and server software need to be upgraded on a somewhat regular basis, but untold numbers of workstations also need attention.

Depending on the size of your organization, your Microsoft Exchange server might be the most robust server in the closet, and finding a suitable replacement for Exchange is quite often a show-stopper. Citadel is a groupware solution that allows organizations not only to avoid upgrading software, but it also runs on a significantly lower-powered machine, thus breaking the hardware upgrade cycle for years to come.

How to Get It

It's always good practice to install and test everything on a test server before moving it into a production environment. Swapping out your mail server is certainly no different, and you should keep your Citadel testing as far away from your production system as possible. Obtaining and configuring Citadel is several orders of magnitude easier with an Internet-connected server, however, because you can avail yourself of the Easy Installation process.

As of this writing, the most current version of Citadel is 6.84. I highly recommend a trip by the Citadel site in order to obtain the most current version of the server and the most current version of the installation instructions. Our testing environment consisted of Debian Sarge with a 2.6 kernel running in VMware Player 1.0.2. For no particular reason, we selected the Web server installation option, but virtually any installation category should work, as Citadel installs everything it needs. In the past, we have installed and run Citadel on a Debian Sarge server proper, and in both cases the installation was flawless.

Installation

I cover the Easy Installation method here not only because it's easy, but it's also fairly undemanding on resources and therefore quite likely that anyone can make use of it. Just about the only requirement for the Easy Installation method is a working—and preferably fast—Internet connection.

The Easy Installation method requires the toolchain, or build environment, to be present on the target platform. In addition, curl (or wget) is required. If you'd like to support SSL connections to the server, you also need libssl-dev. On a Debian system, use the following command to install or verify your build environment:

```
apt-get install build-essential curl libssl-dev
```

Before installing, it is worth noting that Citadel is designed as a black box system running on your server. Part of that black box means that Citadel authenticates logins against its own user database and not against your system user database (typically /etc/passwd). If you'd like Citadel to authenticate against your system user database, you must export the IS_AUTOLOGIN variable to the environment prior to running the Citadel install, like so:

```
export IS_AUTOLOGIN=yes
```

Now that the environment is set, it's time to kick off the Easy Install with the command:

```
curl http://easyinstall.citadel.org/install | sh
```

or, if you'd prefer to use wget:

```
wget -q -O - http://easyinstall.citadel.org/install | sh
```

Citadel downloads, unpacks and starts the installation process. You need to pay attention to the installation process, as Citadel asks all the right questions, but you won't need any of your arcane configuration logs to answer them.

Citadel is humble, and although it brings a lot of power to the party, it doesn't assume that you want

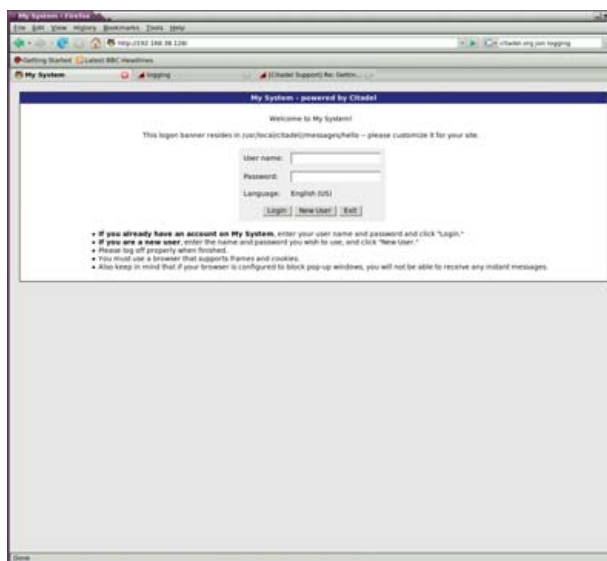


Figure 1. The Login Screen




Figure 2. The Lobby. When Logged In as Administrator

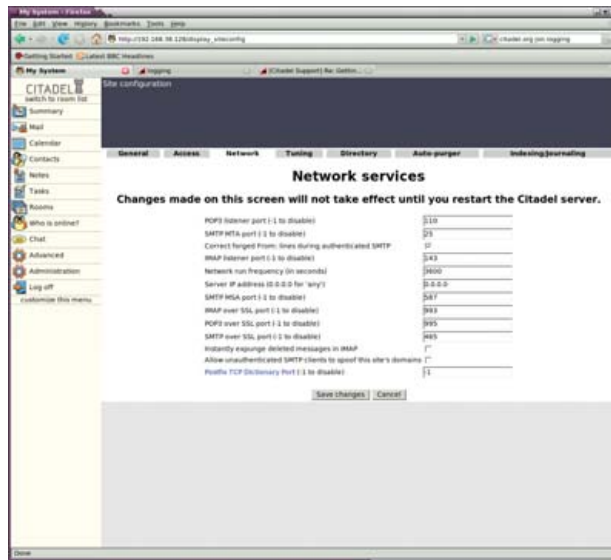


Figure 3. Network Services Settings

any of it. Citadel will ask if you want to use its built-in POP, SMTP or IMAP servers or leave any of your own up and running.

Further, there is a Web interface, called WebCit, which users can make use of to get all of their e-mail, calendar and contact information when on the road or otherwise away from their local e-mail and Personal Information Manager client. If you elect to install WebCit, Citadel won't assume that you want it running on port 80. It therefore is possible to run WebCit on a nonstandard port and leave any existing Web sites you have on port 80 undisturbed.

For the curious, Citadel is installed to `/usr/local/citadel`, and WebCit, if chosen for installation, is installed in `/usr/local/webcit`. Supporting libraries can be found in `/usr/local/ctdlsupport`.

Uninstallation

Uninstalling a Citadel instance installed via the Easy Install method is easy:

1. Delete the three directories mentioned above (`/usr/local/webcit`, `/usr/local/citadel` and `/usr/local/ctdlsupport`).
2. Remove the Citadel and WebCit entries from the `inittab` file (typically `/etc/inittab`).
3. Type the command `init q` to restart init.

Gone.

Initial Configuration

We used the WebCit Web interface to configure and use our Citadel server, but underneath the nice GUI beats the heart of a text-mode BBS. Virtually all of the configuration and much of the daily use of the Citadel system can be used via the text mode Citadel client à la the BBS scene of days gone by. Sadly, that method of communication is largely lost to most modern-day users, so we focus only on WebCit to get the job done.

Having said that, we still need to log in to our Linux server for other reasons, so we have to change the way that Citadel logs. By default, Citadel logs to the console, and that needs to be redirected somewhere else in order to get any work done. There are a variety of different ways to do this, but since Linux provides a configurable `syslog` daemon, it seems logical to edit the `/etc/syslog.conf` file (on Debian) and point the `local4` facility to a log file or somewhere else out of the way.

The first person to log in to the new Citadel Web interface becomes the administrator-level user. To create the administrator account, point your Web browser to the host and port where you told WebCit to listen during the installation, enter a user name and password, and press the New User button (Figure 1). You'll know you've become the administrator if you see the Administration button on the bottom left of the menu when you're logged in (Figure 2).

To enter the site-wide configuration, click the Administration button, and you'll be brought into a well-organized and complete settings menu. Main categories are along the top of the page, and clicking each one brings up the settings for that particular area. As mentioned, Citadel is also a text-mode BBS underneath the WebCit interface, and some of the configuration options make that quite obvious.

Although a good study of all of the configuration items is outside the scope of this article, the most important settings are under the Network and Directory (if you're using LDAP). Under the Network tab, you can enable and disable services and modify the ports on which they run. Under the Directory tab, you can specify your LDAP settings. If you're not using LDAP, you probably can leave both of these screens alone, because the Network defaults either are quite reasonable or will reflect your installation choices (Figure 3).

You may want to take a quick trip to the Access tab and make sure it reflects how you would like new users to sign up. Likely, for a corporate server, administrators will create all of the accounts and user-driven account creation can be shut off.

Sometime before letting users at the WebCit interface, you'll likely want to customize the site a bit. As you navigate

Citadel is humble, and although it brings a lot of power to the party, it doesn't assume that you want any of it.

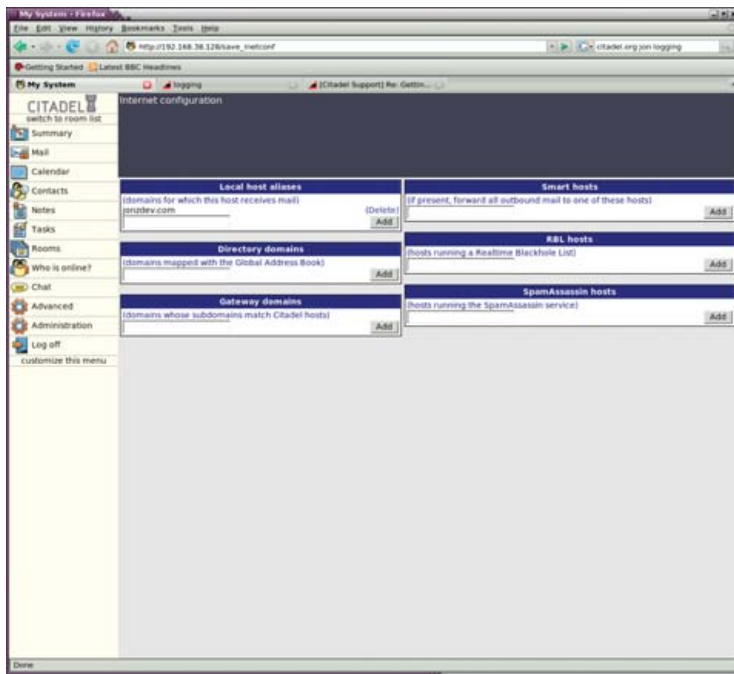


Figure 4. Internet Configuration

through the default WebCit installation, you may notice default text banners on the site that contain the path to their locations. A good example of this is the "Welcome to My System" banner on the main WebCit log in page (Figure 1). A variety of text files exist in `/usr/local/citadel/messages` that can be tailored to your needs.

Setting Up E-Mail

First things first, and before you point your mail records to your new Citadel server, you have to tell it what domains to accept e-mail for. I much prefer Citadel's way of handling this as opposed to mucking about in configuration files. To specify the domains for which you're interested in handling e-mail, click on the Advanced menu option, and then the Domain names and Internet e-mail configuration link.

In the resulting page, enter the first domain for which you want to accept mail in the Local host aliases field. Click the Add button, and continue entering more domains as your situation requires (Figure 4).

The Local host aliases field is the only setting that absolutely has to be filled out, but you may want to integrate some more-advanced functionality within this screen as well. You can specify the domains to map to the Global Address List (GAL), indicate smart host addresses if your server isn't sending mail directly or point to a SpamAssassin or real-time blackhole list (RBL) host to scrub incoming mail before it's delivered.



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

**Hardware Systems For The
Open Source Community—Since 1989**
(Linux, FreeBSD, NetBSD, OpenBSD, Solaris, MS, etc.)

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

AMD Opteron™ Value Server- \$795

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



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

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



1U SCSI Quad AMD Opteron™- (Rev.F) Starting @ \$5293.20

1 of 4 AMD Opteron 8212 CPU
12 GB DDR2 ECC REG PC-3200
1 of 3 73GB SCSI Drive
2 GigaE, CD, FD,
Optional Remote Management Card (IPMI)
Call for more choices of AMD socket F Servers.



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

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



Your Custom Appliance Solution

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



Custom Server, Storage, Cluster, etc. Solutions

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



2354 Calle Del Mundo, Santa Clara, CA 95054

www.asacomputers.com

Email: sales@asacomputers.com

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

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

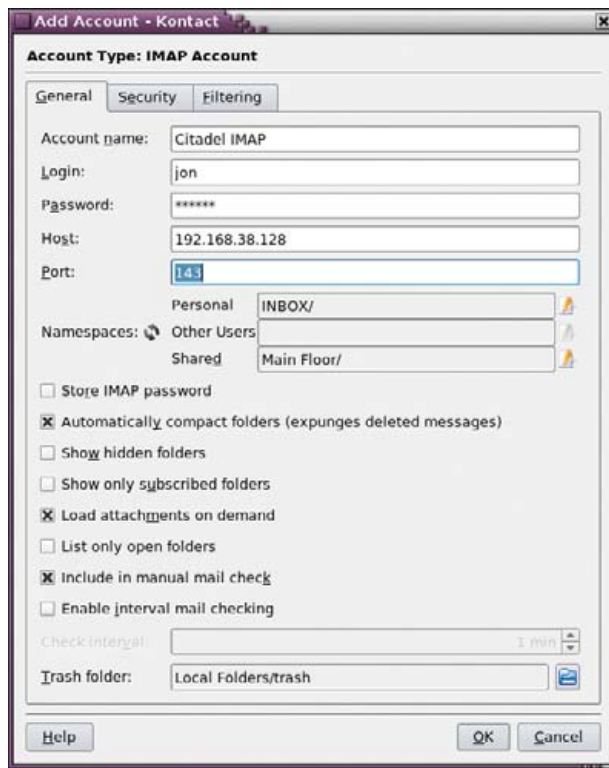


Figure 5. Set Up KMail for Citadel

That's it. You now can send and receive e-mail out of your Citadel installation.

Setting Up Clients

There's no technical reason why a local client has to be set up at all. WebCit exposes all of the most-used groupware functionality via a Web interface, and users can begin using that immediately to organize their lives. However, local clients do bring some power to the table, and many users won't be satisfied with a Web interface. Therefore, onward we go.

Depending on the needs of your users, a variety of Linux clients can replace Microsoft Outlook. After many setups, we've found that KDE's Kontact is the easiest personal information manager to back onto a Citadel server, so that's what we use here.

Kontact is the KDE Project's all-in-one personal information manager. In a sense, Kontact simply provides a unified interface to access KMail, KOrganizer, KAddressbook and some notes and news components.

KMail

Setting up KMail is a rather intuitive process. If you've ever set up a mail client before, you'll be able to set up KMail without issue. As long as you've set up at least one of your Citadel server's IMAP or POP servers, you can set up KMail to use either. Simply plunk in the URL or IP of your Citadel server, your account credentials and be done with it (Figure 5).

KOrganizer

Setting up the calendaring functionality of Kontact is a little more indepth. We've found that the GroupDAV protocol is

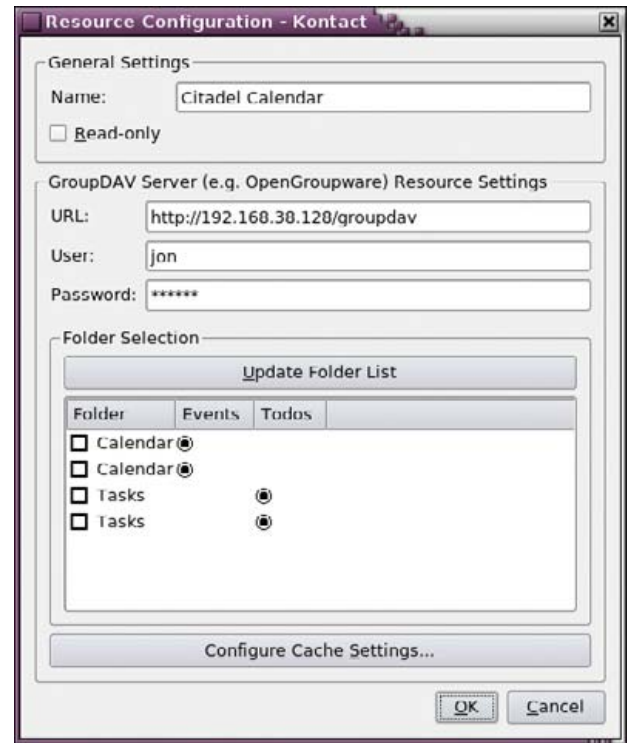


Figure 6. Set up KOrganizer (Kontact) calendar for Citadel.

the easiest and most powerful to set up, so that's what we do here.

One of the few things you need to know is how to construct your GroupDAV URL. Quite simply, your GroupDAV URL is the URL to your Citadel server (including the nonstandard HTTP port if you've told Citadel to listen on a port other than 80) with /groupdav appended to it. In my case, my GroupDAV URL is http://192.168.38.128/groupdav.

To enable KCalendar's groupware functionality, click on the Calendar icon in the left-side pane. At the bottom of the middle pane is a section labeled Calendar. Right-click anywhere in that pane, and select Add. In the resulting window, select the GroupDAV Server option. If you don't see the GroupDAV Server option, it's likely you don't have the kdepim-kresources package installed. Install it, restart Kontact, and you should be good to go.

The Resource Configuration window opens. Enter a name that means something to you in the Name field and your special GroupDAV URL into the URL field. Your user and password credentials are the same ones that you set up when you logged in to Citadel the first time. Click the Update Folder List button, and the bottom Folder Selection pane should populate with Calendar and Tasks radio buttons (Figure 6).

It seems that clicking the check boxes beside the Calendar and Tasks items would enable those items, but the system is a little buggy. In many cases, two instances of Calendar and Tasks show up, as shown in Figure 6. Further, to enable a Calendar or Tasks item, the only way that seems to work is to right-click each item and select Enable from the context menu.

Once you've enabled the Calendar, you can enter items either within Kontact or within WebCit, and the

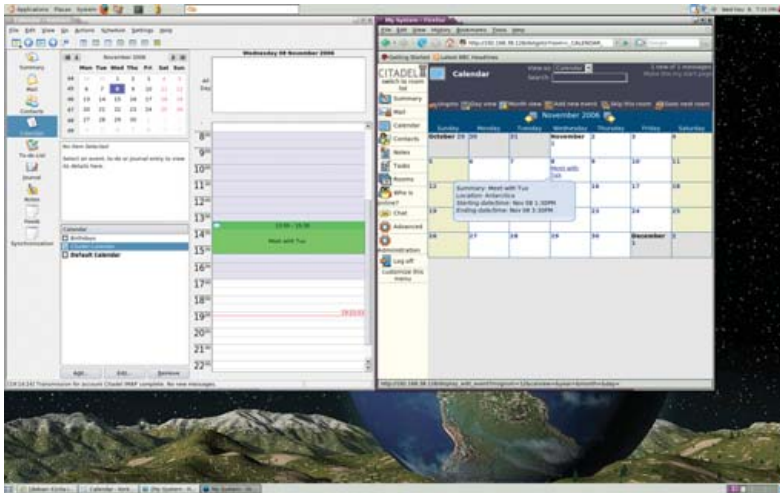


Figure 7. Kontact in Action

items synchronize as mail is checked or other server contact occurs (Figure 7).

Contacts

Setting up Kontact's Contacts (say that five times fast) is much the same as setting up KCalendar. Click the Contacts icon in the left-side pane. At the bottom of the middle column is a pane labeled Address Books. The right-click trick doesn't work here, so click the Add button instead. Select the same GroupDAV Server option, and fill in all the same data that you filled in for the KCalendar setup. Click the Refresh Folder List button, use the right-click-to-enable trick, and you're off to the races (Figure 8).

As with KCalendar, once you've set up your GroupDAV

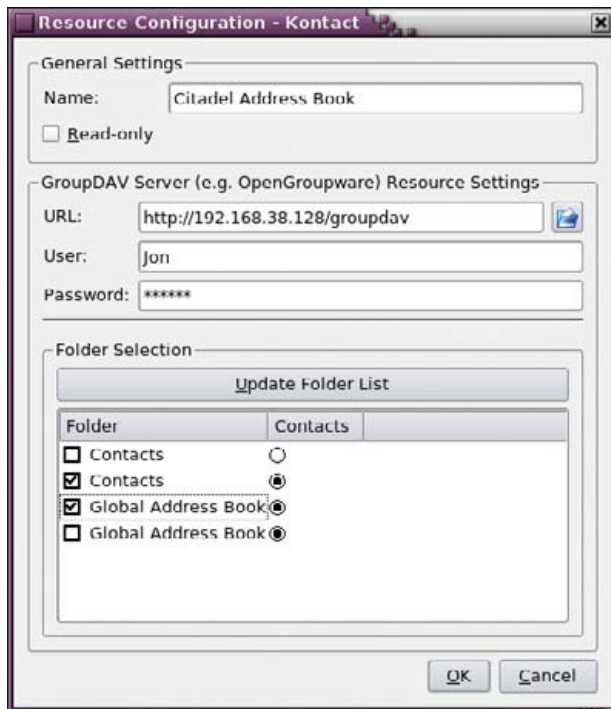


Figure 8. GroupDAV Settings in Kontact

connector, you can now manage your contact data from either KDE or WebCit (Figure 9).

Tasks and the Journal are just plain-old work once KCalendar is set up. They don't require any of their own setup.

A lot of other clients support the GroupDAV protocol to varying degrees. Any of these can be used in place of Kontact, albeit likely with less functionality. For a complete list of clients and the status of their GroupDAV support, go to the GroupDAV site (www.groupdav.org/implementations.html).

GroupDAV isn't the only technology that can be used with Citadel. WebDAV and Webcal can be used with clients, such as Mozilla Sunbird and Evolution, to share calendars and schedule events. There is also a Microsoft Outlook connector in the works, but at the moment, Outlook can be used to access only POP/IMAP e-mail

and IMAP folders. As time marches on, more and more clients that support GroupDAV and WebDAV come onto the scene. The Citadel FAQ contains a maintained list of clients and how to configure them.

Although a few groupware projects are underway that can give Microsoft Exchange a run for its money, we've found that Citadel is quite simply the easiest to install and maintain. The hardest part of a Citadel install is waiting for all the components to download. Citadel is under active development, and by the time this article prints, a new version may be out. The lead developer, Art Cancro, can be found in the Citadel support on the UNCENSORED! BBS forums (uncensored.citadel.org), along with other Citadel developers and experienced users. ■

Jon Watson (www.jonwatson.ca) is a Canadian GNU/Linux enthusiast who regularly contributes articles to the Linux community. When not writing, blogging and podcasting about free and open-source software, Jon frequently can be found in his office polishing his Linux+ certification, which impresses no one but himself.

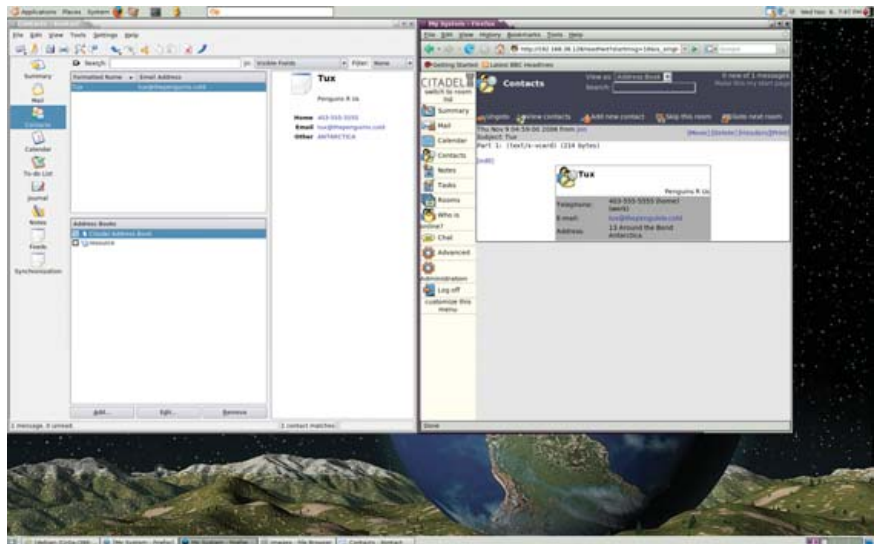


Figure 9. Citadel and Kontact Accessing the Same Data

Interview with Christof Wittig and Jerry Fiddler of db4objects

db4objects emerges as a unique blend of company and community. NICHOLAS PETRELEY

LJ: Can you tell me a little about the history of db4o?

Christof: db4o is the native Java and .NET object database developed since 2000 by a small group of developers around Carl Rosenberger. Like many open-source projects, db4o was driven by users who felt there was an urgent need for an object persistence solution that was more efficient and better performing than incumbent solutions based on relational paradigms or flat file/serialization.

In 2004, after the product was successfully delivered to a handful of early customers, we felt it was ready for prime time. I was asked to join the company as CEO to grow the business. We named the company db4objects. The company is based in Silicon Valley.

db4objects brought db4o to the masses by adopting the open-source/dual-license model, as we know it from MySQL, and by raising funds from Silicon Valley luminaries, such as Mark Leslie, founding CEO of Veritas (and our chairman); Vinod Khosla, founding CEO of Sun Microsystems; and Jerry Fiddler, founding CEO of Wind River, who joined our board earlier this year. Jerry is here with me today to add his insight regarding db4objects.

LJ: What is db4o, and how are most of your customers using it?

Christof: db4o is an embeddable object persistence solution—and because this is such an awkward term, we call it an object database. However, databases often are associated only with Oracle, Versant, AS/400—large, single-unit, server-side solutions, and with a DBA. db4o is an embeddable component that helps to persist objects in distributed, multi-unit, client-side applications, where no DBA is present—that's why the term database is correct, though a bit ambiguous here.

Target applications span a broad range: devices (such as photo-copiers and consumer electronics), mobile applications (such as doorstep delivery systems), packaged software (on a PC) and server-side middleware, and caches (such as in a SCADA system for railways or pipelines).

The unifying aspect is zero administration, reliability, high performance even with highly constrained resources, and it has to be a Java and .NET environment.

LJ: Would it be accurate to say that the market for db4o exists due to the popular adoption of a JVM in embedded systems?

Jerry: Yes. As Christof said, db4o goes where Java and .NET go. Rather than “embedded systems”, I much prefer to talk about “device software”.

Traditionally, most device software was written in C or C++,

often with a strong focus on hardware rather than best software practices. But devices are becoming ever-more sophisticated, with more memory, faster CPUs and, maybe most important, rich connectivity. Both the software and the data they deal with are becoming much more complex. The software needs to ratchet up a notch, and Java and .Net provide two ways to do that. The RYO (roll-your-own) approach doesn't fly anymore—it's too cumbersome, too slow, not agile and doesn't connect well.

In many industries and application types, but especially in mobile, you now see more and more practices and technologies from servers and PCs being adopted on smaller devices, thanks to Moore's Law.

It's this convergence of technologies and practices, where Java embedded (and .NET with Windows mobile) is gaining ground rapidly. And, db4o is the only optimized persistence solution database for this space. This is one reason why I am so excited about this company.

LJ: What are the key differences in version 6?

Christof: Performance, performance, performance.

LJ: That's it?

Christof: That's most of it, yes. Any product in this space should focus on performance, because it's the number-one concern of developers. A product that doesn't address performance in this space either is not user-driven or simply doesn't have any.

For the new db4o version 6, it's not only about being a little faster. Version 6 has up to ten times faster response times, in some cases even 250 times. So it's not about being a little faster, it is about enabling a whole new set of functions (such as complex, ORed queries) and applications (for example, for small devices, where we have decreased the demand on memory consumption up to 90% and are making it more deterministic).

LJ: How did all these highly significant changes come about?

Christof: There's a technical and there's an organizational answer to this.

The technical answer is the use of a new B-Tree index architecture, first introduced in v5.5 a few months ago, which we could leverage to achieve these amazing results.

The organizational answer is what links all of this to open source. Our user community has driven the product road map for this release. Their clear number-one priority is performance, but there also are many other, less spectacular but relevant improvements, such as faster defragmentation, a new server-side cursor technology and an improved, less Java-esque .NET API.

LJ: So how did you orchestrate version 6 in a user-driven fashion?

Christof: We started off in early 2006 with a user survey where we gathered more than 1,000 responses and votes. We then gathered with some users at the first db4o User Conference in July 2006 in London, where we spec'd out some of the details. And while we were producing, we were getting feedback, sometimes in two-hour intervals—because that's the release cycle of our continuous builds! Even our weekly planning meetings are now Skype-casted, so the core developer team always remains in touch with the community of users out there.

LJ: How much of the changes in version 6 were the result of commercial developer input vs. the Open Source community.

Christof: We don't make this distinction. Even the contracted and paid developers are as much a part of the community as any community member is part of our company. There is not an “us” and “them”; there is only a “we”. It's a collaborative approach. You earn your say by merit (ideas and contributions), not by rank, seniority, title or politics.



Christof Wittig

It's this convergence of technologies and practices, where Java embedded (and .NET with Windows mobile) is gaining ground rapidly.



Jerry Fiddler

At our board meetings, we feel this invisible stakeholder (the user community) sitting right there—and she has a big say in what we do and how we allocate resources.

LJ: We've been hearing database developers cry for a good database without the overhead of SQL for a long time. I know comparing something like MySQL or PostgreSQL to a JVM-based database is largely apples to oranges, so I'm sorry if this question seems unfair. But, why do you think you're among the first, if not the first, to listen to this demand?

Christof: It's the focus on the developer rather than on the DBA. We don't go well with DBAs—and its enterprise IT organization. We go well with applications where the developer "owns" the data in terms of anticipating its administration needs in the application rather than relying on runtime database administration. Just think of the software in your car. You expect your vendor to take care of the database administration, don't you? Or, do you want to hire a DBA to defrag your central car CPU?

There are many applications out there that the (relational) database vendors don't see (because they talk to CIOs and DBAs day in and day out). So, although they all have "light", "compact" or "embedded" products, they are usually just teaser products to lure you into their full-blown server version. But these products aren't optimized for embedded, nor are their sales organizations. As a result, none of the big vendors holds more than 5–10% market share, while 50% of the developers, according to Chris Lanfear's research at VDC, still write their own persistence solutions today!

Jerry: Combine this with the power of open source, a totally new paradigm, which is essentially a low-cost production and distribution model for software, and then you can see that a small, focused player like db4objects (just like Sleepycat, JBoss or MySQL) can have a real impact in the market place, despite its initially more limited resources.

Because db4o is so profoundly connected to the user community, developers know that the users like and want the software. Because of the dual-license model, developers can work with the software immediately, without going through management and purchasing, at least for the prototyping stage. Then, developers can go to management with a working, demonstrable project. Now, there's really no decision left for the CIO to make—it's a done deal. It turns the entire decision-making process upside down. From a db4o perspective, it enables very low sales and marketing costs, which in turn feeds back and enables this very democratic process of acquiring users and customers. Compared to a traditional software sales and marketing structure, it's really lovely.

So the timing is right—from the market, the open-source perspective and the technology perspective.

Jerry: That's the really exciting part of db4o. At Wind River, we were always working hard to find ways to connect and engage the developer community, but it was always a challenge. The Open Source world, and especially db4o, has really honed the connection of the user community into the company in a seamless and low-friction manner.

The way db4o has built and nurtured its community of now 15,000 registered developers is amazing. I think it's the only way to build a company in this space today, and db4objects has done a fantastic job to walk their talk. Open-source and community collaboration is not only a label, it's also a true commitment. The user community is an integral part of product development and support.

LJ: Do you have an SQL compatibility layer for those who want to use db4o but want to use it in the same way they might use another SQL database?

Christof: Sure. But it's not a layer, it's a replication service called dRS (db4o Replication System).

Why would you want to access your data on your smartphone with SQL? You access it only with your application.

But what is even better is to sync this data with back-office systems, which typically store data in the tables of Oracle or MySQL. Then, you access it with your report writer or data-mining tool right there where you really need it. The dRS, powered by Hibernate, provides this bilateral synchronization with all major RDBMSes, making db4o entirely compatible with legacy systems.

This is an excellent example, by the way, of how the community drove our product decisions. We contemplated internally, like *all* closed-source object database vendors before us, to provide an SQL access layer, because people wanted to check this item off their list in database evaluations. Sales teams often pressure you to make decisions that are not necessarily the best—"we must have SQL or we don't win this deal."

When we proposed the idea, our users begged us not to do it. They did not want us to contort db4o for the sake of some managers

who simply don't get it. They said, "Make it data-compatible with the dRS, but don't bring the entire overhead of SQL back through the back door. We don't want it!"

Can you imagine? It was a rebellion of the developers against their managers, who forced them to use tools that simply didn't give them the ability to do their jobs.

Previously, 50% of our target developers ended up writing their own database! SQL wasn't good, but anything other than SQL wouldn't pass their managers' desks for approval.

And then along came open source. Open source empowered developers to make autonomous decisions. They can now go to their managers with a proof of concept and show: "See, we can ship 10% faster than our competition, and this is due to the fact that I disrespected the SQL-only policy in our organization." This is where real innovation comes from!

LJ: There seem to be similarities in your business model to things like Qt and MySQL, both of which are also dual-licensed commercial and open-source projects. I get the sense, however, that there are key differences too. Can you elaborate on those differences?

Christof: I am very familiar with MySQL since I am a Stanford researcher and have worked extensively with it.

MySQL has a relational database, and, by definition, is not particularly suited for embedded use. Most relational databases, based on SQL—a DBA language—are built for and sold to end users (DBAs and their CIOs). RDBMSes are not doing too well as OEM products, and, in fact, the MySQL Network subscription, an end-user offering, currently is overtaking MySQL's traditional embedded business in revenue.

Our product, based on native Java or .NET—the developer's language—is sold to product developing companies as an OEM component. Being open source, it's attractive for evaluation and optimization. For redistribution in non-GPL'd products, people then choose the commercial license without the GPL constraints, but with indemnification and specific support packages.

LJ: Can you give me a Java code example of how someone might open a database table, query the table and fetch some fields?

Christof: Wow. I'm just the CEO...Jerry can you help?

OK, here's something with respect to queries, showing how we cater our product to Java or .NET developers, not DBAs.

Here's a query in plain Java with what we call "Native Queries" for all students under 20 and with grade A:

```
List<Student> students = database.query<Student>(new Predicate(){
    public boolean match(Student student){
        return student.getAge() < 20
            && student.getGrade().equals(gradeA);}});
```

And here's the equivalent in JDOQL, to take one example for SQL:

```
String param = "String gradeParam"
String filter = "age < 20 & grade == gradeParam";
Query q = persistenceManager.createQuery(Student.class, filter);
q.declareParameters(param);
Collection students = (Collection)q.execute(gradeA);
```

What do you see? Pure Java in the first, strings in the second. The strings are aliens. They are like a little bit of a Chinese poem in an English poem—they simply don't fit.

What's the benefit of being native to Java? I'll give you two reasons, though there are many more:

1. What is typesafe—do you know that age is a number field? With db4o you'll know. The IDE will give you a type-mismatch while you write. With JDOQL, you won't know until runtime. So, you are simply more productive as a developer with a native solution.
2. How to refactor—if you need to change "age" to "_age" with Java/db4o it's one simple step with your IDE, for example Eclipse. With JDOQL (and any other like JDBC, such as EJBQL, HQL and you name them), you're stuck, you need to do it manually (or not refactor at all). You don't want to make a find/replace on the source code, do you?

Now, not refactoring means increased bugs, more maintenance costs, less reuse of software, as we all know. So it comes down to developer productivity once again.

With db4o, you build better and leaner code faster. You can go home at five.

Don't get me wrong: it's not a solution for everything, for example, to build an application that connects to a legacy database. But if it is appropriate, such as in a mobile app, where there's no legacy, you shouldn't really break your fingers and induce all the SQL complexity that is of no use at all, as there's no DBA around.

Of course, you still can use SQL for compliance purposes, but then you may find that your competitor has it faster, more agile and leaner, less buggy, more performance and more feature-rich—with db4o. That's not where you want to be as a developer or as a company.

LJ: Can you mention some examples of companies that have adopted db4o? How are they using the product, and how has that productivity played out?

Christof: Boeing builds the P-8A Multi-Mission Maritime Aircraft for the US Navy with db4o. Boeing says, "db4o provides the advantages of significantly lower database administration and improved developer productivity. db4o helps Boeing to manage development costs and schedules while also reducing operational costs."

Version 6 has up to ten times faster response times, in some cases even 250 times.

Ricoh in Tokyo with \$17 billion in sales just decided to build its future copier models with db4o. Testuo Ito, the software lead, says, "db4o provides a persistence solution for our broad range of technical challenges and for our stringent quality standards. After a long period of evaluation, we found that db4o has the flexibility to fit our cutting-edge architectures, which aim to achieve better productivity in our object-oriented software development."

BOSCH Sigpack, worldwide leader in fully automated packaging technology with \$800 million in sales, relies on db4o to deliver its Delta XR-31 robot. "Our biggest concern is shortening our commissioning time. The use of db4o on the data back end has helped us to achieve a time-saving effect of at least 10% on each project."

Intel says, "db4o will make application development much easier for our group. The OR mapper/SQL database alternative really did not allow us to do everything we needed and forced us to contort our application designs. By comparison, implementing with db4o was seamless and worked within our existing architecture."

Want more?

LJ: Thanks, that is impressive. But how does that scale? What kind of adoption rate has db4o seen over the past few years?

Christof: We now have 15,000 registered developers and potentially 2–3x unregistered developers—in some two years, nearing 1,000,000 downloads. So far, we have closed 200 commercial design win deals. We're growing more than 100% each year.

LJ: That's quite a rapid rate of increase in adoption. Given that the new version has such a dramatic performance increase, how do you expect that to play in your future growth? I would expect performance to be a huge factor in embedded systems, so I would expect a big increase. But is there any other factor I'm missing that could add to or subtract from a big spike in adoption?

Christof: Performance is crucial, but there's a wide spectrum.

We will, for instance, provide better performance to our client/server users. Not that you run db4o as a client/server application on a mobile phone, but there are many more instances where multiple (embedded) clients connect to a (embedded) database server. It's wild what's going on out there. That should make us, again, attractive for a larger and larger crowd.

LJ: Because much of your market consists of embedded systems, I would assume most of your registered users are using db4o in devices they sell. That means they have to use the commercial version, right? If my assumption is wrong, why?

Christof: No.

We define embedded—like Carl Olofson from IDC—different from devices. Embedded to us means "invisible to the end user". And that's a much larger market than device software, though with similar characteristics. Packaged software is a good example for embedded, yet not device software. And strictly speaking, the mobile applications that our ISV customers sell on Handango for the smartphone are not "device software", but still embedded. All these are our core market.

In our community we also have end users, that use (sometimes "abuse" it for prototyping)—small systems, nonprofits, academia. That's great, so we get more people to love db4o, but sometimes they go beyond the scope of what db4o is designed for.

The commercial part kicks in when someone wants to redistribute db4o and does not want to go GPL with its own intellectual property. These guys then come to us and look for an alternate, safer licensing option, which we provide with our flexible OEM commercial terms. They also look for indemnification and a direct support option,

when someone responds in 24 hours or less, guaranteed.

Of course, we also have embedding users under the GPL, such as ITAnyplace, an open-source mobile platform that extends applications and content to mobile devices.

Any user counts and is welcome, as she strengthens our ecosystem and is lost revenue to our closed-source competitors.

LJ: What other products or add-ons do you have available for db4o?

Christof: I mentioned the db4o Replication System (dRS), to provide compatibility with back-end RDBMSes, such as Oracle and MySQL.

The other add-on is the ObjectManager, a UI that comes in a total rewrite dubbed v6.0, including many user demands for the handling of large data sets and lots of enhancements around database inspection and application debugging.

These tools are open source too, and available under dual licenses.

LJ: Thanks again for taking the time to talk with us!

Christof and Jerry: Thanks for your interest and your time! ■

Nicholas Petreley is Editor in Chief of *Linux Journal* and a former programmer, teacher, analyst and consultant who has been working with and writing about Linux for more than ten years.

LINUX ONLINE TRAINING
Real Training on Real Servers

Linux Administration Course - **\$295.00**
* CentOS or Ubuntu Track

Firewall Design Course - **\$195.00**

Desktop Course - **\$34.95**
* Ubuntu, SUSE, Fedora, or DSL

SpiderTools.com

Virtual Filesystems Are Virtual Office Documents

Use **libferris**, **XML** and **XSLT** to create virtual filesystems and virtual documents. **BEN MARTIN**

Virtual filesystems can be made into writable virtual office documents. The old UNIX slogan “everything is a file” together with the `xslfts://` virtual filesystem allows for transparently editing relational databases, RDF and arbitrary XML with OpenOffice.org.

The **libferris** virtual filesystem presents both files and their metadata as a virtual filesystem. The boundaries of what is considered a filesystem by **libferris** include such interesting data sources as PostgreSQL, LDAP and Firefox as well as standard Web items, such as HTTP, FTP and RDF.

Many virtual filesystems allow directory contents to be synthesized from other directories. The classic example of this is a union filesystem where a collection of existing filesystems are taken as input to generate a filesystem showing the set union of the base filesystems.

Recently, the **libferris** filesystem has gained support for performing XSLT on a filesystem and exposing the result as a virtual filesystem. To keep things simple, I refer to the original virtual filesystem as the input filesystem and the filesystem that results from the XSL transform as the translated filesystem. As the main use of XSL is to describe translations on trees, this fits nicely for the use of creating translated filesystems.

Although there are differences between a **libferris** filesystem and the XML data model, there are also many similarities. A file’s contents map to the text content of an XML element. A file’s metadata is exposed by **libferris** as Extended Attributes (EAs), which map to XML attributes on the file’s XML element. A notable difference between a filesystem and an XML data model is that the document ordering in XML is not always easy to preserve. To keep the mapping simple, a file can generate only one text node in an XML document. Technically, an XML element can have multiple text nodes as children.

Because of the close relation with the XML data model, the **libferris** filesystem supports viewing any filesystem as a Document Object Model (DOM), which is created on demand. The inverse also is true: you can expose a DOM as a filesystem. As **libferris** can mount XML as a filesystem, the lines between what is a filesystem and what is XML are somewhat blurred.

Many modern applications store their documents as XML files. As filesystems and XML can be interchanged with **libferris**, this allows you to use those applications to edit filesystems directly. The main problem with having such applications edit filesystems directly instead of XML is that the schema of the application’s XML file usually does not match the layout of the filesystem.

This is where `xslfts://` can be used to create a translated filesystem that matches the layout the application is expecting. For example, you could take a table in a PostgreSQL database as the input filesystem and have the XSL message that table into a virtual spreadsheet file, which you load into OpenOffice.org.

The possibilities become even more interesting when write support in the translated filesystem is considered. After you have made some changes to the above virtual spreadsheet file in

OpenOffice.org, you “save” the file. The filesystem then applies a reverse XSLT and updates the input filesystem (in this case a PostgreSQL table) to reflect your changes.

To support this, you have to have two XSL files. The first stylesheet translates an input filesystem into the format you are interested in. The second XSL file (the reverse stylesheet) provides the inverse translation. In the future, the second XSL file should become optional if it can be inferred from the actions of the initial translation.

Reverse stylesheets can specify updates either using explicit URLs for each file to change or relative paths. The explicit URLs method expects the reverse stylesheet to specify the absolute URL for each file to be updated. This can be convenient for `xslfts://` applications where URLs play a role in both the source and translated filesystem. For example, when editing some RDF files with OpenOffice.org, the subject URI will be available to allow the reverse stylesheet to use explicit updates.

The relative paths method is conceptually similar to applying diff and patch to your filesystems. The reverse stylesheet generates a list of changes to make using a relative path for each file to change. Some options from the patch utility are available to the reverse stylesheet as well. The root element can contain a strip attribute that works similarly to the strip option of patch. The autocreate attribute, when set to true, will make **libferris** try to create new files where the reverse stylesheet specifies a relative path that does not exist in the source filesystem.

Currently, both reverse stylesheets must supply the entire contents of each file to update. This is not a major drawback, as that information already will be fully available in the translated filesystem.

The following sections show two uses: creating new virtual filesystems and directly interacting with them from the console and creating virtual office documents. This is followed by some advice for creating custom stylesheets by hand.

Manufacturing Filesystems with `xslfts://`


Translated filesystems can be accessed through the `xslfts://` scheme. This filesystem can be interacted with using the **libferris** clients or exposed using Filesystem in Userspace (FUSE) through the Linux kernel.

As **libferris** allows you to see an XML file as a filesystem, the XML file shown in Listing 1 will be used as the input filesystem.

The XSL file shown in Listing 2 will create our translated filesystem from the input filesystem. It is important to keep in mind that although the input filesystem in this case is generated from an XML file, it could just as easily be data from a mounted LDAP server. The XSL will create two elements under the document root element. The `file3` element will have the original contents of the virtual “file” for `file3` in the input filesystem. The `file7` will have the attribute `myattr` as its contents.

The translated filesystem can be used just like any other filesystem with the command-line utilities `ferrisls`, `fcat`, `ferriscp` and so on. The `xslfts://` URL scheme in **libferris** lives above most other URL schemes

SOFTWARE DEVELOPMENT WEST 2007 CONFERENCE & EXPO
MARCH 19-23, 2007 | SANTA CLARA CONVENTION CENTER | SANTA CLARA, CA



**EARLY BIRD
DISCOUNT**
Register by FEB 23
**SAVE UP
TO \$300**

20
CELEBRATING
YEARS!

OVER 200
SESSIONS IN
14 IN-DEPTH
TRACKS

- Business of Software
- C++
- Java
- .NET
- Ruby
- Modeling & Design
- People, Projects & Methods
- Requirements & Analysis
- Security
- Testing & Quality
- Web 2.0
- Web Services/SOA
- XML

**STROUSTRUP
& SUTTER ON C++
BACK BY POPULAR DEMAND:**

Join C++ creator Bjarne Stroustrup and renowned C++ expert Herb Sutter for an in-depth two-day tutorial on C++.

NEW THIS YEAR: RUBY, WINDOWS VISTA AND WEB 2.0!

PLUS: Keynotes, Expo, Birds-of-a-Feathers, Panels, Case Studies, Roundtables, Parties and Special Events

REGISTER TODAY AT WWW.SDEXPO.COM

SD
WEST 2007

Listing 1. example.xml

```
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<root>
  <file1 size="200"/>
  <file3>filesystems inside XML?</file3>
  <file7 myattr="foo" >Something blue</file7>
</root>
```

Listing 2. example.xsl

```
<?xml version="1.0" encoding="utf-8"?>
<xsl:stylesheet
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  version="1.0"
  >
  <xsl:output method="xml"/>

  <xsl:template match="/">
    <root>
      <xsl:apply-templates/>
    </root>
  </xsl:template>

  <xsl:template match="file3">
    <context original-url="{@url}" name="file3">
      <xsl:value-of select="@content"/>
    </context>
  </xsl:template>

  <xsl:template match="file7">
    <context original-url="{@url}" name="file7">
      <xsl:value-of select="@myattr"/>
    </context>
  </xsl:template>
</xsl:stylesheet>
```

and allows you to materialize a filesystem at any point by supplying an XSL transform to apply. The location of the XSL files themselves is determined based on an `xsltfs` path you set in `libferris`. The use of an `xsltfs` path avoids embedding full stylesheet paths into `xsltfs://` URLs. As the stylesheets are specified using a CGI-like syntax, avoiding the use of the `/` character means that there is no ambiguity for filenames in `xsltfs://`.

You can apply a stylesheet at any point in your virtual filesystem. The result of applying a stylesheet to the `example.xml` filesystem will become the contents of a directory rooted at the `example.xml?stylesheet=example.xml` virtual directory.

Without any use of `/` in the `xsltfs://` parameters, the filename and parameters together are used to specify the name of a virtual directory that `xsltfs://` makes on demand. Because there is no unambiguity, you then can navigate directly into the translated filesystem rooted at this virtual directory. This is shown in the examples below.

Part of a filesystem is shown in Listing 3 to make things clearer. I have applied the `foo.xml` to the `example.xml` file using the special CGI-like syntax to name a virtual directory. `libferris` creates this virtual directory for me to allow direct navigation into the translated filesystem. The `rootElement` is the root of the translated filesystem; in XML terms, it is the document root of the result of applying the `foo.xml` stylesheet

Listing 3. Generating a Translated Filesystem for example.xml

```
xsltfs://
context
file
tmp
  example.xml
  example.xml?stylesheet=foo.xml
  rootElement
    myFoo1
    myBar2
```

Listing 4. Exploring Our New XSLT Filesystem

```
$ bash
$ URL='xsltfs://context/file/tmp/example/
➔example.xml?stylesheet=example.xml'
$ cd /tmp/example
$ ls
example-rev.xml  example.xml  example.xml
$ export LIBFERRIS_XSLTFS_SHEETS_URL='pwd'
$ ferrisls -l $URL
  0 root
$ ferrisls -l $URL/root
 23 file3
  3 file7

$ fcat $URL
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<root>

  <context name="file3"
  original-url="file:///tmp/example/example.xml/root/file3"
  >filesystems inside XML?</context>

  <context name="file7"
  original-url="file:///tmp/example/example.xml/root/file7"
  >foo</context>

</root>

$ fcat $URL/root/file3
filesystems inside XML?
```

to the filesystem rooted at `example.xml`. Filesystems live inside the context subdirectory of `xsltfs://` to allow other parameters and expansion to be done in `xsltfs://` at a later time.

The `xsltfs` path can be set using the XSLT stylesheets page of the `ferris-capplet-general` configuration tool. In addition to setting the XSLT path with `ferris-capplet-general`, you can use the `LIBFERRIS_XSLTFS_SHEETS_URL` environment variable to pass in the path explicitly where your forward and reverse stylesheets are located. This makes using `xsltfs` with the FUSE module from shell scripts quite simple, as you do not need to install your stylesheet files. Stylesheets can be stored in any filesystem `libferris` can reach.

For the purposes of this example, I have the files and stylesheets stored in `file:///tmp/example`. If I am running my examples from the `example` directory, it is sufficient to put `.` into my XSLT path—see the example in Listing 4.

Listing 5. example-rev.xsl

```
<?xml version="1.0" encoding="utf-8"?>
<xsl:stylesheet
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  version="1.0"
  xmlns:ferris="http://libferris.org"
  exclude-result-prefixes="ferris"
  >
  <xsl:output method="xml"/>

  <xsl:template match="/">
    <explicit-updates>
      <xsl:apply-templates/>
    </explicit-updates>
  </xsl:template>

  <xsl:template match="context[@name='file3']">
    <context url="{@original-url}">
      <xsl:value-of select="."/>
    </context>
  </xsl:template>

  <xsl:template match="context[@name='file7']">
    <attribute url="{@original-url}"
  name="myattr"><xsl:value-of select="."/></attribute>
  </xsl:template>

</xsl:stylesheet>
```

Things become more interesting when we provide a reverse stylesheet, as shown in Listing 5. In this case, we are mapping things back fairly plainly to where they originated in the input filesystem. The file7 content is placed back into the myattr XML attribute of the input document. Having an explicit reverse XSL transform provides you with the freedom to update only part of the original filesystem as you see fit. You also can use functions from the stylesheet to modify the data on its way back to the input filesystem.

Now that we have the forward and reverse XSL, we can happily modify the contents of the original example.xml file by interacting with the virtual file(s) in our xslfs://, as shown in Listing 6.

The example in Listing 6 shows two options for updating your filesystem: either by changing individual virtual files or by updating the virtual XML document (the translated filesystem) in a single shot. The first method of updating individual files maintains the filesystem metaphor in the xsltfs. The second method of updating via rewriting the main virtual XML document provides support for XML editing applications, such as OpenOffice.org where a document is read, manipulated and rewritten.

The URLs can be quite ugly and rather long. If you are editing such filesystems frequently, you might want to expose the xsltfs using FUSE. Editing virtual XML files with OpenOffice.org requires the use of FUSE to expose the virtual XML file through the Linux kernel.

Virtual Office Documents

If the format of the output of xsltfs:// is well known, such as an OpenOffice.org document, you can create file format automatically from the XSL files.

The ferris-filesystem-to-xsltfs-sheets client is used to set up stylesheets automatically. A plugin system is used to allow new file



ASA COMPUTERS

Want your business to be more productive?

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

Hardware Systems For The Open Source Community—Since 1989

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

Dempsey/Woodcrest Server Starts at \$2395

- 1U Dual Core 5030 CPUs.
- 4GB FBDIMM Memory.
- Supports upto 64GB FBDIMM.
- 120 GB hotswap hard drive.
- 2xIntegrated Dual 10/1000 LAN.



Dual Xeon Server starts at \$4139

- 5U Dual Xeon 2.8 Ghz 800 FSB.
- iSCSI or NAS Software options.
- 8x120 GB Sata Hard drives -Upto 18TB.
- 512 MB RAM
- Fail hard drive LED indicator.



Dual Xeon 800FSB Storage starts at \$6,699


- 8U Dual Xeon 2.8 Ghz 800FSB CPUs.
- 2TB of storage (36TB max).
- 1GB RAM
- NAS or iSCSI software options
- 2 x 10/100/1000 Gigabit LAN.



Your Custom Appliance Solution

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

"Your Logo Here"



ASA Collocation

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

ASA Collocation Special

First month of collocation free.*

Storage Solutions

IDE, SCSI, Fiber Raid solutions
NAS, DAS, iSCSI, SATA, SAS
3Ware, Promise, Adaptec
JMR, Kingston/Storage solutions

Clusters

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

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



2354 Calle Del Mundo,
Santa Clara, CA 95054

www.asacomputers.com

Email: sales@asacomputers.com

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



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

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

Listing 6. Changing the XML File through our New XSLT Filesystem

```
$ bash
$ URL='xsltfs://context/file/tmp/example/
example.xml?stylesheet=example.xsl&
reverse-stylesheet=example-rev.xsl'
# Change the file3 element to have new content
$ echo foo | ferris-redirect -T $URL/root/file3
$ cat example.xml
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<root>
  <file1 size="200"/>
  <file3>foo
</file3>
  <file7 myattr="foo">Something blue</file7>
</root>

# Update everything based on a new XML file
$ cat example-update1.xml
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<root>
  <context name="file3" original-url=
  "file:///tmp/example/example.xml/root/file3"
  >A new file3 text node
  </context>
  <context name="file7" original-url=
  "file:///tmp/example/example.xml/root/file7"
  >A new file7 myattr</context>
</root>

$ cat example-update1.xml | ferris-redirect -T $URL

$ cat example.xml
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<root>
  <file1 size="200"/>
  <file3>A new file3 text node
  </file3>
  <file7 myattr="A new file7 myattr"
  >Something blue</file7>
</root>
```

Listing 7. Allowing a User to Use FUSE on Fedora Core

```
root-bash-# usermod -a -G fuse ben
```

formats to be supported in the future. To see which plugins are available, use the `--plugin=help` command-line option.

You need to use a FUSE filesystem in order to read and write virtual office documents directly. This also can be set up automatically by the `ferris-filesystem-to-xsltfs-sheets` client using the `--fuse=foo` command-line option.

Some distributions require additional setup for a user in order to use FUSE mounts. On Fedora Core, you have to add the user to the `fuse` group, which can be done as shown in Listing 7.

An example of setting up a little PostgreSQL table and creating a new virtual office document to allow editing this table is shown in Listing 8.

The final command in Listing 8 opens the virtual spreadsheet document, which should look similar to Figure 1. I then changed

Listing 8. Setting Up a Virtual Office Document to Edit a Database Table

```
bash-$ psql
ben=# create database lj;
ben=# \c lj;
You are now connected to database "lj".
lj=# create table msgs
lj-# ( id serial primary key,
lj-#   num int, msg varchar(200),
lj-#   foo varchar(100) );
lj=# insert into msgs values
lj-# ( default, 7, 'This is msg #1', 'Foo is Bar');
lj=# insert into msgs values
lj-# ( default, 12, 'Second message', 'ii tenki');
lj=# select * from msgs;
 id | num |      msg      |   foo
----+-----+-----+-----
  1 |  7 | This is msg #1 | Foo is Bar
  2 | 12 | Second message | ii tenki
(2 rows)
\q

bash-$ ferrisls pg://localhost/lj
msgs
bash-$ ferrisls --xml pg://localhost/lj/msgs
<ferrisls>
<ferrisls url="pg://localhost/lj/msgs" name="msgs">
  <context id="1" num="7"
    msg="This is msg #1" foo="Foo is Bar"
    name="1" primary-key="id" />
  <context id="2" num="12"
    msg="Second message" foo="ii tenki"
    name="2" primary-key="id" />
</ferrisls>
</ferrisls>

bash-$ ferris-filesystem-to-xsltfs-sheets \
--plugin excel2003 --fuse msgs \
pg://localhost/lj/msgs

bash-$ ferrisls -lh ~/ferrisfuse
... ben ben 129 06 Oct 21 11:56 mount-msgs.sh
... ben ben 4.0k 06 Oct 21 11:56 msgs

bash-$ cd ~/ferrisfuse/
bash-$ ./mount-msgs.sh
bash-$ ls -lh msgs
... 0 ben ben 3.8K Jan  1 1970 msgs.xml*
bash-$ cat msgs/msgs.xml | head
<?xml version="1.0" encoding="UTF-8" ... ?>
<Workbook xmlns=...>
  <OfficeDocumentSettings xmlns=...>
    <Colors>
      ...

bash-$ ooffice msgs/msgs.xml
```

some data in the second row and saved the file giving the result shown in Figure 2.

Looking at the PostgreSQL table after saving the virtual office document shows the updated contents—see Listing 9.

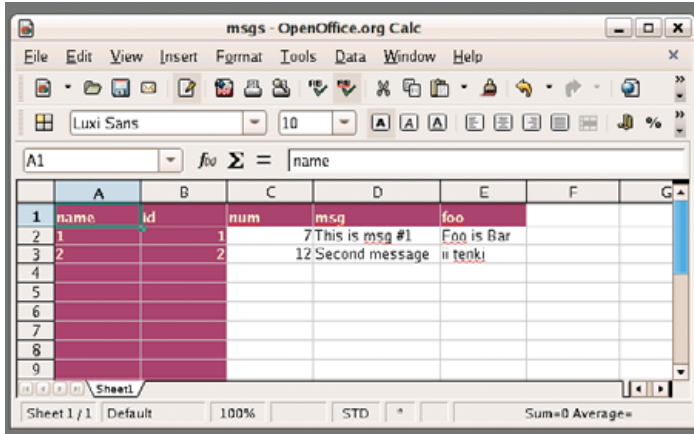


Figure 1. Initial View of Virtual Office Document

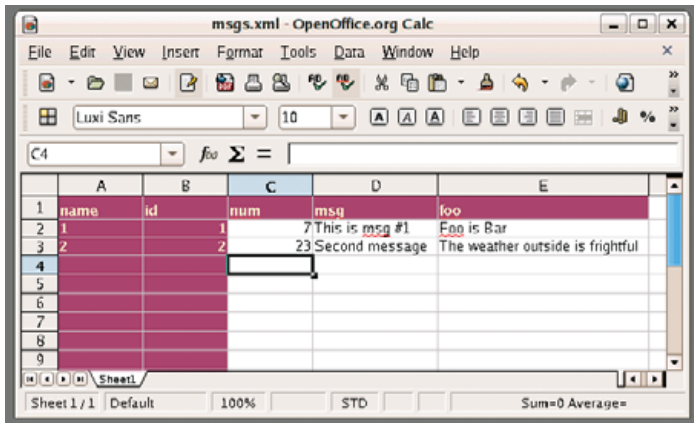


Figure 2. Some changes to the second row are saved back to the database.

Listing 9. The Contents of the Database after Editing with OpenOffice.org

```
bash-$ psql lj
lj=# select * from msgs;
 id | num | msg | foo
-----+-----+-----+-----
  1 |  7 | This is msg #1 | Foo is Bar
  2 | 23 | Second message | The weather outside...
(2 rows)
```

Google Earth and xsltfs://

The ferris-mount-etagere-as-kml.sh script uses xsltfs:// and FUSE to set up a read/write virtual KML file. The stylesheets translate between libferris geoemblems and the KML format for place names used by Google Earth.

The stylesheets used to expose libferris emblems provide an example of translating a whole tree in libferris into a hierarchical XML document for an external application to use. The is-dir EA from the input filesystem is used to determine the type of XML element to generate in the translated filesystem, as KML files require the use of Placemark or Folder elements depending on whether children are to be found.



Great Minds Great Solutions

1-877-25-SERVER www.genstor.com

Customized Solutions for

SERVERS :: STORAGE :: APPLIANCES

Linux - FreeBSD - x86 Solaris - MS etc.

SERVERS

LOW POWER - BUDGET - HIGH DENSITY

(1U, 2U, 3U and above - Xeon, P4, Opteron, VIA)



2U Dual Xeon/Opteron Server
Upto 24GB RAM
Upto 8 SATA/SCSI HDD
Starting @ \$ 2000.00
Call for your custom needs

SERVER MANAGEMENT & MORE

Genstor offers Data Center-proven Linux Management Technology - which gets you :

Linux Server Management Made Easy

- Provisioning
- Deployment
- Change Management for Linux Environments large and small
- Track Changes
- Instant Rollback and much more

Setup your racks with heterogeneous Linux environments in minutes
Contact sales@genstor.com

STORAGE

SATA- NAS-DAS-iSCSI -SAS Storage Solutions



5U Dual Xeon/
Opteron SATA Storage
Upto 24GB RAM
Upto 26 SATA HDD
Upto 18TB Storage
Starting @ \$ 3930.00
Call for your custom needs

CUSTOM APPLIANCE SOLUTIONS

Prototype - Certifications - Logo Screen Printing



Custom Turnkey OEM Appliance Solutions
From Prototype to Drop Ship
Custom OS/Software Image Installs
No Quantity is small to Customize
Call for your Custom OEM Solutions

FACING POWER PROBLEMS AT DATA CENTER PL. CALL FOR SOLUTIONS

Contact Genstor For Your Hardware Needs

Genstor specializes in customizing hardware around the OS of your choice (Linux, *BSD, x86 Solaris etc.) with Intel and AMD offerings. Please contact Genstor sales@genstor.com for all your hardware needs.

GENSTOR SYSTEMS, INC.

780 Montague Exp. # 604, San Jose, CA 95131

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

Email: sales@genstor.com

Fax: 1-408-383-0121

Prices, Products and Availability subject to change without notice.

Listing 10. Developing and Debugging New Stylesheets

```

$ ferrisls -R --xml-xsltfs-debug \
  --show-ea=name,content,myattr \
  example.xml/root

<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<ferrisls>

  <root name="root"
    url="file:///tmp/KK/example.xml/root">

    <file1 content="" myattr=" " name="file1"
      url="file:///tmp/KK/example.xml/root/file1"/>
    <file3 content="filesystems inside XML?"
      myattr=" " name="file3"
      url="file:///tmp/KK/example.xml/root/file3"/>
    <file7 content="Something blue" myattr="foo"
      name="file7"
      url="file:///tmp/KK/example.xml/root/file7"/>
  </root>

</ferrisls>

$ ferrisls -R --xml-xsltfs-debug \
  --show-ea=name,content,myattr \
  example.xml/root >| input.xml

$ FerrisXalanTransform -s example.xml -m input.xml
transform XML:input.xml with xsl:example.xml
<?xml version="1.0" encoding="UTF-8"?><root>
  <context
    original-url="file:///tmp/KK/example.xml/root/file3"
    name="file3">filesystems inside XML?
  </context>
  <context
    original-url="file:///tmp/KK/example.xml/root/file7"
    name="file7">foo
  </context>
</root>

</context>
</root>

$ export LIBFERRIS_XSLTFS_SHEETS_URL=`pwd`
$ URL=xsltfs://context/file/tmp/example/example.xml/
  ↳root?stylesheet=example.xml
$ fcat $URL
<?xml version="1.0" encoding="UTF-8" standalone="no" ?>
<root>
  <context name="file3" original-url="file:///home/ben/xsltfs/
  ↳example.xml/root/file3">filesystems inside XML?
  </context>
  <context name="file7" original-url="file:///home/ben/
  ↳xsltfs/example.xml/root/file7">foo
  </context>
</root>
$ fcat $URL >| translated.xml
$ vi translated.xml
..make changes to test reverse sheet
..inserting CHANGE_A and changeB into the elements

$ FerrisXalanTransform -s example-rev.xml \
  -m translated.xml

transform XML:translated.xml with xsl:example-rev.xml
<?xml version="1.0" encoding="UTF-8"?>
<explicit-updates>
  <context
    url="file:///home/ben/xsltfs/example.xml/root/file3"
    >filesystems inside CHANGE_A XML?
  </context>
  <attribute
    url="file:///home/ben/xsltfs/example.xml/root/file7"
    name="myattr">foo changeB
  </attribute>
</explicit-updates>

```

Rolling Custom Stylesheets

For testing purposes, if the `LIBFERRIS_XSLTFS_DONT_UPDATE` environment variable is set, `libferris` performs the reverse stylesheet application and logs what updates would have been done but does not actually update the input filesystem.

There are a few hints that can make setting up and adjusting custom forward and reverse stylesheets much simpler.

I use the `example.xml` file shown in Listing 1 again here as the input filesystem. Although in this example, I am starting with `example.xml`, which is an XML file, we want to see how `libferris` sees this input filesystem, not only the raw XML itself. For example, the contents of an elements text nodes will be available as the content attribute when `libferris` mounts this XML file.

To get at the `libferris` view of the XML, I use `ferrisls` with its `--xml-xsltfs-debug` option. I also need to recurse the `example.xml` file to get the whole filesystem and explicitly select any attributes that the `example.xml` file will want to use.

The manual application of a forward stylesheet is shown in Listing 10.

The reverse stylesheet can be applied to the translated filesystem XML file. Once this output looks sane, non-destructive testing can be done by applying it through `xsltfs://` with `LIBFERRIS_XSLTFS_DONT_UPDATE` set. Make sure `ferris-logging-xsltfs` is set to debug in the `ferris-capplet-`

logging configuration tool to get all the information about what would have been updated.

Some Future Directions

The major planned feature is the automatic derivation of the reverse stylesheet. This would make setting up `xsltfs://` mountpoints much simpler. Things, such as duplicating nodes in the forward XSL file, would require an explicit reverse XSL file to resolve conflicts where each duplicate was edited in the transformed filesystem.

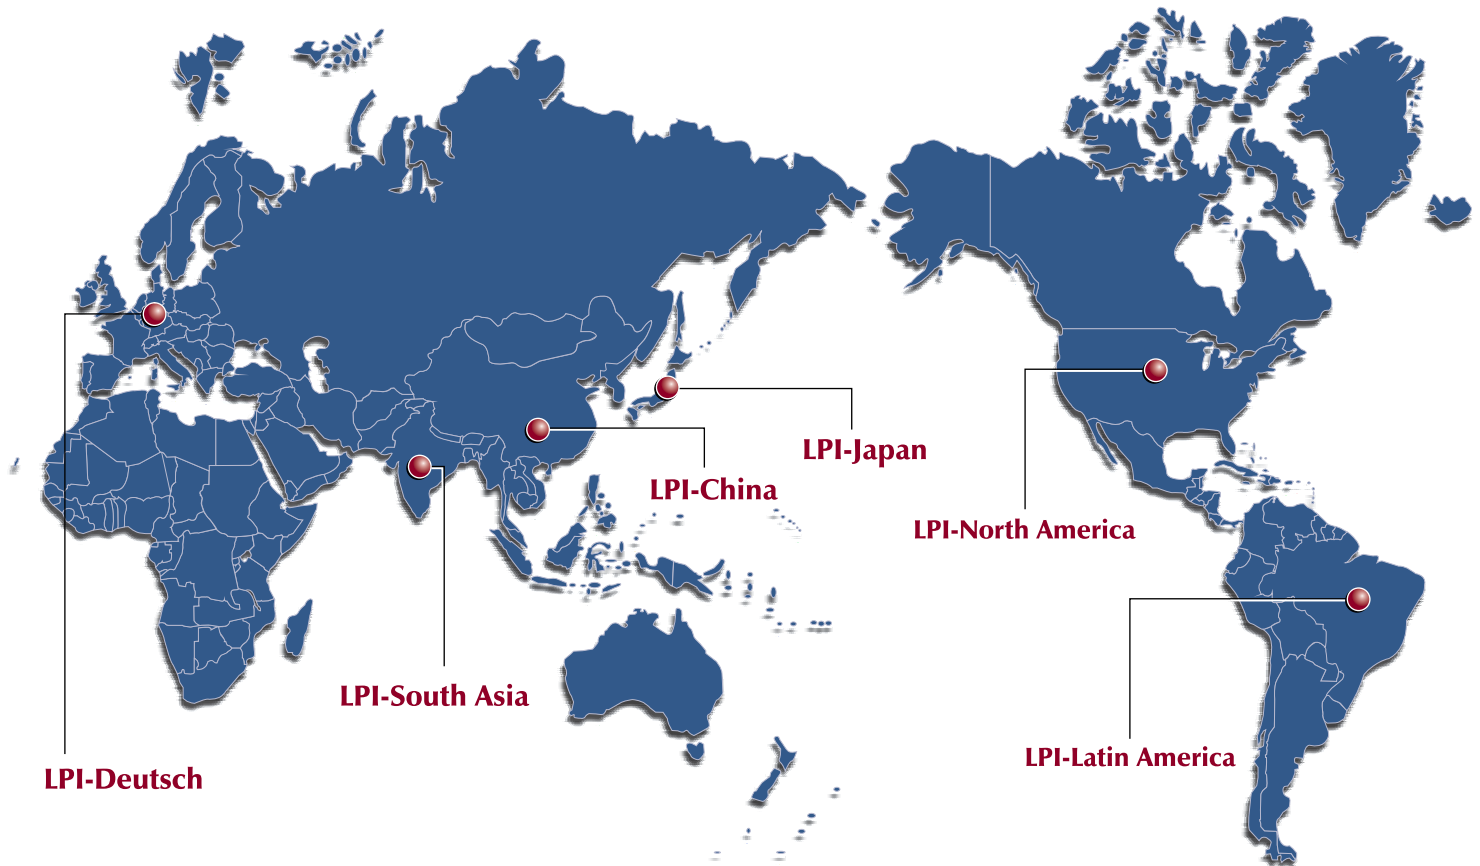
More plugins for `ferris-filesystem-to-xsltfs-sheets` are in the cards. For example, being able to edit data from common LDAP schemas, such as user authentication in OpenOffice.org, would be nice. Support for creating virtual OpenOffice.org zip files as the target of `xsltfs://` would allow the creation of native OpenOffice.org documents.

More of the command-line options of patch probably will become available for the reverse stylesheet to use. ■

Resources for this article: www.linuxjournal.com/article/9513.

Ben Martin has been working on filesystems for more than ten years. He is currently working toward a PhD at the University of Wollongong, Australia, combining Semantic Filesystems with Formal Concept Analysis to improve human-filesystem interaction.

Growing a World of Linux Professionals



We at the Linux Professional Institute believe the best way to spread the adoption of Linux and Open Source software is to grow a world wide supply of talented, qualified and accredited IT professionals.

We realize the importance of providing a global standard of measurement. To assist in this effort, we are launching a Regional Enablement Initiative to ensure we understand, nurture and support the needs of the enterprise, governments, educational institutions and individual contributors around the globe.

We can only achieve this through a network of local "on the ground" partner organizations. Partners who know the sector and understand the needs of the IT work force. Through this active policy of Regional Enablement we are seeking local partners and assisting them in their efforts to promote Linux and Open Source professionalism.

We encourage you to contact our new regional partners listed above.

Together we are growing a world of Linux Professionals.



Stable. Innovative. Growing.

Simple Access Berkeley DB Using STLdb4

STLdb4 makes C++ programming with the Berkeley DB simpler and more effective.

BEN MARTIN

The **Berkeley DB library** provides a solid implementation of both the B-Tree and Hash file structures. The implementation includes support for transactions, concurrent access of database files from multiple processes, and secondary indexing as well as logging and recovery.

In this article, I use the term database to refer to a B-Tree or Hash maintained by Berkeley DB. These databases allow rapid key to value look-ups.

The standard distribution of Berkeley DB comes with both a C and C++ API. Unfortunately, the standard Berkeley DB C++ API is a very thin wrapper neglecting modern C++ designs, such as smart pointers, standard C++ I/O streams, iterators, default arguments, operator overloading and so on. As a concrete example of the lack of reference counted smart pointers, the Berkeley DB API for `Db::get()`, shown in Listing 1, includes two `Dbt` pointers and the ownership of the memory for these is not immediately obvious.

Listing 1. Standard Berkeley DB C++ API `Db::get()`

```
#include <db_cxx.h>
int Db::get(DbTxn *txnid, Dbt *key, Dbt *data,
           u_int32_t flags);
```

The STLdb4 Project was created to make using the Berkeley DB from C++ easier. The STLdb4 API aims to make simple database interaction trivial while still keeping more advanced usage simple. A Berkeley DB object behaves similarly to an STL collection allowing look-ups and the setting of elements using an overloaded array operator. A full example program is shown in Listing 2. After execution, the file named with `argv[1]` will contain a Berkeley DB B-Tree file containing the foo-bar data pair.

The main class is the `Database` and the reference counted smart pointer for this class is called `fh_database`. This trend is used throughout STLdb4 where the smart pointer for `Foo` is called `fh_foo`. Databases can be opened either as in Listing 2 directly in the constructor or using the empty constructor and the `open()` or `create()` methods later. The main difference between `open` and `create` is that `create` requires a database type (for example B-Tree or Hash) and will create a new database at the given path if none exists already.

In the example in Listing 2, I don't have to close the database explicitly, because the smart pointer to the `Database` object will handle that for me.

Standard STL collection methods, such as `empty()`, `size()`, `insert()`, `erase()`, `count()`, `begin()`, `end()`, `find()`, `upper_bound()` and `lower_bound()`, all exist in the `Database` class. There are also partial versions of the latter three methods. The partial versions allow the looking up of entries with part of a key in B-Tree files. A bidirectional iterator object is returned by many of the above methods.

When storing large values in the database, using the standard I/O

Listing 2. STLdb4 Setting and Getting Values

```
#include <iostream>
#include <STLdb4/stldb4.hh>

using namespace STLdb4;
using namespace std;

int main( int argc, char** argv )
{
    fh_database db = new Database(DB_BTREE, argv[1]);
    db["foo"] = "bar";
    cerr << "foo is set to:" << db["foo"] << endl;
    return 0;
}
```

streams can be more efficient than using the `get()` method or overloaded array operator. This is because the standard I/O streams use partial read and write operations on the underlying Berkeley DB file. A standard I/O stream is obtained using the `getIStream()` and `getOStream()` methods of the `Database` class.

The example in Listing 3 shows the standard C++ I/O stream interface for STLdb4. The housekeeping of performing partial I/O to the Berkeley DB file is handled by STLdb4. Accessing large chunks of data through this API maintains a low memory consumption. The API shows one of the used `getOStream()` calls as having a `ferris_ios` first parameter. As the `libferrisstreams` library that STLdb4 uses offers generic I/O stream support, the `ferris_ios` is a backward-compatible extension of the `std::ios` bitfield. The extension allows specifying such things as memory mapped backing and sequential stream access to be nominated for use where supported. The output from running this example is shown in Listing 4.

Storing Objects

One major difference between the `Database` class and an STL collection like `std::map<>` is that the key and value are not parameterized in `Database`. The main reason for this is that the items in a `Database` object are usually not in RAM but are read from disk on demand. Also, in order not to limit the functionality offered by Berkeley DB, the `Database` class has to support storing arbitrary data and not a heterogeneous collection of objects.

The illusion of stored objects can be created using implicit constructors and type conversion thin object wrappers. Shown in Listing 5, the `Person` class stores some information about people. The implicit constructor takes a `DatabaseMutableValueRef`, which is the class returned by the array operator in `Database`. A `Person` object is implicitly convertible to an `std::string` to enable it to be serialized to disk. As the main

Listing 3. Standard C++ I/O Streams for Berkeley DB Files

```
#include <iostream>
#include <STLdb4/stldb4.hh>

using namespace STLdb4;
using namespace Ferris;
using namespace std;

int main( int, char** )
{
    fh_database db = new Database( DB_BTREE,
                                "/tmp/play.db" );

    string data = "1234567890";
    db[ "fred" ] = data;
    cerr << "Initial value:" << db["fred"] << endl;

    {
        fh_iostream ss = db->getIOStream( "fred" );
        ss << "54321";
    }
    cerr << "Second value:" << db["fred"] << endl;

    {
        fh_iostream ss = db->getIOStream( "fred" );
        ss.seekp( 3 );
        ss << "AAAA";
    }
    cerr << "post seekp value:" << db["fred"] << endl;

    // truncate the iostream and write
    {
        Database::iterator di = db->find( "fred" );
        fh_iostream oss = di.getIOStream( ios::trunc, 0 );
        oss << "sm";
    }
    cerr << "Trunc and write:" << db["fred"] << endl;

    // append some more data to end of iostream
    {
        fh_iostream oss = db->find( "fred" )
            .getIOStream( ios::ate, 0 );
        oss << "AndMore";
    }
    cerr << "at end write value:"
        << db["fred"] << endl;

    return 0;
}
```

Listing 4. Output of Standard C++ I/O Streams Example

```
Initial value:1234567890
Second value:5432167890
post seekp value:543AAAA890
Trunc and write:sm
at end write value:smAndMore
```

Advertiser Index

For advertising information, please contact our sales department at 1-713-344-1956 ext. 2 or ads@linuxjournal.com.
www.linuxjournal.com/advertising

Advertiser	Page #	Advertiser	Page #
ABERDEEN, LLC www.aberdeeninc.com	21	LPI www.lpi.org	79
ACMA www.acma.com	87	MBX www.mbx.com	13
AML www.amltd.com	41	MICROWAY, INC. www.microway.com	C4, 17
APPRO HPC SOLUTIONS appro.com	C2	NOWTHOR CORPORATION www.nowthor.com	95
APTITUDE CORPORATION www.sarcheck.com	31	OPEN SOURCE SYSTEMS, INC. www.opensourcesystems.com	5
ASA COMPUTERS www.asacomputers.com	67, 75	PENGUIN COMPUTING www.penguincomputing.com	29
AUTHENTUM www.authentium.com/linuxjournal	88	POGO LINUX www.pogolinux.com	33
AVOCENT CORPORATION www.avocent.com/ice	1	POLYWELL COMPUTERS, INC. www.polywell.com	61
CARL.NET www.cari.net	55	THE PORTLAND GROUP www.pgroup.com	3
CORAD, INC. www.coraid.com	37	RACKSPACE MANAGED HOSTING www.rackspace.com	C3
CYBERMONTANA, INC. www.spidertools.com	71	R CUBED TECHNOLOGIES www.rcubedtech.com	9
EMAC, INC. www.emacinc.com	83	SD WEST www.sdexpo.com	73
EMPERORLINUX www.emperorlinux.com	15, 57	SERVERS DIRECT www.serversdirect.com	27
FAIRCOM CORPORATION www.faircom.com	19	SILICON MECHANICS www.siliconmechanics.com	25, 39
FREE SOFTWARE FOUNDATION member.fsf.org	94	SUPERMICRO www.supermicro.com	35
GECAD TECHNOLOGIES/AXIGEN www.axigen.com	85	SWELL SOFTWARE, INC. www.swellssoftware.com	50
GENSTOR SYSTEMS, INC. www.genstor.com	77	TECHNOLOGIC SYSTEMS www.embeddedx86.com	91
HURRICANE ELECTRIC www.he.net	23	TYAN COMPUTER USA www.tyan.com	7
IRON SYSTEMS www.ironsystems.com	34	VERIO www.verio.com	11
LINUX JOURNAL www.linuxjournal.com	6, 43	WILEY TECHNOLOGY PUBLISHING www.wiley.com	93
LINUX CERTIFIED, INC. www.linuxcertified.com	45		

function shows, this thin wrapper makes it appear that the Database is storing Person objects.

Listing 5. Storing and Reading Objects with STLdb4

```
#include <iostream>
#include <STLdb4/stldb4.hh>

using namespace STLdb4;
using namespace std;

class Person
{
public:
    string email;
    string name;
    string phoneNum;
    explicit Person( const string& name,
                    const string& email,
                    const string& ph = "" )
        :
        email( email ), name( name ), phoneNum( ph )
    {}

    x
    Person( const DatabaseMutableValueRef& r )
    {
        stringstream ss;
        ss << (string)r;
        getline( ss, name, '\0' );
        getline( ss, email, '\0' );
        getline( ss, phoneNum, '\0' );
    }
    operator string() const
    {
        stringstream ret;
        ret << name << '\0' << email << '\0'
            << phoneNum << '\0';
        return ret.str();
    }
};

int main( int, char** )
{
    fh_database db = new Database( DB_BTREE,
                                  "/tmp/play.db" );

    db->insert(
        make_pair(
            "alex", Person("Alex", "alex@foo.com")));
    db->insert(
        make_pair(
            "barry", Person("Barry", "barry@bar.com")));

    Person p = db["barry"];
    cerr << "Barry has email address:"
         << p.email << endl;

    return 0;
}
```

Secondary Indexing

Sometimes the information that you are storing has multiple keys by which you would like to be able to find a given item quickly. For example, if you are storing contact information, you want to be able to look up people based on either their name or e-mail address.

You could achieve the above by storing each person's information manually, using the name as the key and maintaining a second database from e-mail address to name. To find a person by e-mail address, you would use the e-mail-keyed database to find the name and then the name database to find the actual information. Maintaining indexes manually like this is highly error-prone, and moreover, the secondary indexes in Berkeley DB can do this housework for you automatically.

The above example can be implemented by having the primary key-value data stored with the person's name as the key and a secondary index on the e-mail address(es). This setup is shown in Figure 1. I refer to the database with the name-to-person data mapping as the main database and the e-mail look-up database as the secondary index.




Figure 1. A Secondary Index for Quick Look-Up by E-Mail Address

The main concern when using secondary indexing with STLdb4 is how to extract the secondary key from your data. There are some template functions in STLdb4 to help you with this. The `getOffsetSecIdx()` template takes an offset as its template argument and will return all the data from that offset to the end of an item as the secondary key. The `getOffsetLengthSecIdx()` is similar, but it allows you to specify both the offset and length of the secondary key data. Finally, the `getOffsetNullTerminatedSecIdx()` takes an offset and a string skip count to allow you to extract the *n*th null-terminated string after a given offset. For example, if you have five (32-bit) integer values followed by four null-terminated strings as your persistent format, you could use an offset of 20 and a skip of two to extract the third null-terminated string as your secondary index key.

Assuming the use of the Person class from Listing 5, the code in Listing 6 creates and uses a secondary index on the e-mail address for your Person objects. Because the disk format starts with our string data, when creating the extraction function with `getOffsetNullTerminatedSecIdx()`, I use an offset of zero and skip one null-terminated string (the name) to extract the e-mail address null-terminated string.

I then perform a partial look-up using the secondary index. The `equal_range_partial()` method finds both the lower and upper bound for partial key material. In this case, I find any e-mail addresses that begin with `al`. The output from the program is shown in Listing 7. Note that the first element of the iterator is the key from the secondary index, and the second element is the data from the main database. The key from the main database for this look-up is available through `getPrimaryKey()` on the iterator object.

Transactions

Transactions are supported either by passing an explicit transaction object to each method or by setting the implicit transaction on Database objects. The latter style can be very convenient in cases when the overloaded array operator is used, which does not allow a transaction object to be passed in (only one argument can be passed to the array operator).

Listing 6. Secondary Indexing

```

unlink( "/tmp/play.db" );
unlink( "/tmp/play.sec.db" );

fh_database db = new Database( DB_BTREE,
                             "/tmp/play.db" );
Database::sec_idx_callback f
    = getOffsetNullTerminatedSecIdx<0,1>();
fh_database secdb = Database::makeSecondaryIndex(
    db, f, DB_BTREE, "/tmp/play.sec.db" );

db->insert(
    make_pair(
        "alex", Person("Alex", "alex@foo.com")));
db->insert(
    make_pair(
        "alfred", Person("Alfred", "alfred@bar.com")));
db->insert(
    make_pair(
        "andrew", Person("Andrew", "andy@foo.com")));
db->insert(
    make_pair(
        "barry", Person("Barry", "barry@bar.com")));

pair< Database::iterator, Database::iterator > p
    = secdb->equal_range_partial( (string)"al" );
for( Database::iterator di = p.first;
     di != p.second; ++di )
{
    string prim;
    di.getPrimaryKey( prim );
    cerr << "di... "
         << " primary:" << prim
         << " first:" << di->first
         << " second:" << di->second << endl;
    Person p = di->second;
    cerr << "Person has name:" << p.name
         << " email:" << p.email << endl;
}

```

When explicitly passing a transaction object to the database for each method call, the Transaction class has the commit() and abort() methods either to ensure the data is stored safely on disk or the whole transaction is fully reverted. When the last reference to a transaction object goes out of scope, it will call commit() in its destructor if it was not already committed or aborted.

If you are operating on only one database, you can largely avoid the Transaction class and use the start() method of the Database class to begin an implicit transaction. When using an implicit transaction, the commit() and abort() methods of the Database class perform the transaction finalization actions.

The simplest method of using transactions is shown in Listing 8. Things of note in the example include the use of a database environment, which in this case will include initialization of the Berkeley DB transaction subsystem. A transaction object must be passed to the Database object when it is created. In the Database constructor, I pass a new Transaction that will be handed in as an fh_trans smart pointer, which will clean up the Transaction object for me after the Database object is constructed. When executed, the Initial value and Final value

Listing 7. Secondary Indexing Program Output

```

di...
  primary:alex
  first:alex@foo.com
  second:Alexalex@foo.com
Person has name:Alex email:alex@foo.com
di...
  primary:alfred
  first:alfred@bar.com
  second:Alfredalfred@bar.com
Person has name:Alfred email:alfred@bar.com

```

lines will print the same information to cerr.


The same transaction can be used with multiple databases by holding onto the Transaction object smart pointer and associating it with each database. This is shown in Listing 9. The second part of the example uses setImplicitTransaction() to associate the databases with the current transaction.

The default action of a Transaction object can be changed to calling abort() by setting setDefaultDestructionIsAbort(true) on the Transaction object. This is very handy for use with a Resource

Data Acquisition & Control Computer

iPac 9302

- Cirrus Logic EP9302 ARM9 200 Mhz Processor
- Floating Point Math Engine
- 2 USB 2.0 Host Ports
- SD/MMC Flash Disk Slot
- 40 Digital GPIO Lines
- 1 10/100 Base-T Ethernet port
- 5 channels of 12 bit A/D & 3 PWMs
- 1 RS232 & 1 RS232/422/485 Serial Port
- Battery Backed Real Time clock/calendar
- Eclipse uClinux Development Environment



2.6 Kernel

The iPac has enough I/O for demanding applications & with a size of 3.5" x 3.8" it can fit almost anywhere. Please contact us for more information.

Since 1985
OVER
22
YEARS OF
SINGLE BOARD
SOLUTIONS

EMAC, inc.
EQUIPMENT MONITOR AND CONTROL

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

Listing 8. Using Implicit Transactions with STLdb4

```
#include <iostream>
#include <STLdb4/stldb4.hh>

using namespace STLdb4;
using namespace std;

int main( int,char** )
{
    Environment::setDefault(new Environment( "/tmp" ));

    fh_database db = new Database(
        new Transaction(), DB_BTREE, "/tmp/play.db" );

    db->start();
    db[ "foo" ] = "bar";
    cerr << "Initial value:" << db[ "foo" ] << endl;
    db->commit();

    db->start();
    db[ "foo" ] = "newbar";
    cerr << "Middle value:" << db[ "foo" ] << endl;
    db->abort();

    cerr << "Final value:" << db[ "foo" ] << endl;
    return 0;
}
```

Listing 9. Using Explicit Transactions with STLdb4

```
fh_trans trans = new Transaction();
fh_database db1 = new Database(
    trans, DB_BTREE, "/tmp/play1.db" );
fh_database db2 = new Database(
    trans, DB_BTREE, "/tmp/play2.db" );
db1[ "foo1" ] = "bar1";
db2[ "foo2" ] = "bar2";
trans->commit();

// create a new implicit transaction and go again
trans = new Transaction();
db1->setImplicitTransaction( trans );
db2->setImplicitTransaction( trans );
db1[ "foo1" ] = "bar111";
db2->set( "foo2", "bar222", 0, trans );
// we'd rather not put these changes in after all
trans->abort();
```

Acquisition Is Initialization (RAII) programming style to revert a transaction automatically if any exception occurs in a code block. This RAII style is shown in Listing 10. The code block marked starting at the comment (AA) sets the default destruction action for a transaction to call abort() and then modifies the database with this transaction. An exception is explicitly thrown that will cause the Transaction object to be destroyed (its last reference being the one held by tr on the stack). This will call abort() for the transaction, and we will eventually print the

Listing 10. Using RAII Transactions with STLdb4

```
int main( int,char** )
{
    Environment::setDefault(new Environment( "/tmp" ));

    fh_database db = new Database(
        new Transaction(), DB_BTREE, "/tmp/play.db" );

    try
    {
        {
            fh_trans tr = new Transaction();
            tr->setDefaultDestructionIsAbort( true );
            db->setImplicitTransaction( tr );
            db[ "foo" ] = "bar";
            tr->commit();
            tr = 0;
        }

        // (AA) RAII with transactions
        // Don't use setImplicitTransaction() in this block
        {
            fh_trans tr = new Transaction();

            tr->setDefaultDestructionIsAbort( true );
            db->set( "foo", "First setting", 0, tr );
            Database::iterator diter = db->find("foo",tr);
            diter->second = "this is something evil";
            throw exception();
            tr->commit();
        }
    }
    catch( exception& e )
    {
        cerr << e.what() << endl;
    }

    cerr << db[ "foo" ] << endl;
    return 0;
}
```

old "bar" value at the end of the example.

The Database::iterator class uses Berkeley DB cursors in its implementation, so the transaction we pass to Database::find() will be used for any operations performed on the database iterator. For example, if getIOStream() was called on diter, STLdb4 would be performing partial I/O using the transaction tr on the Berkeley DB file behind the API.

This use of RAII is very handy for code that wants to make changes to the database in one go, but that might throw an exception along the way.

The setImplicitTransaction() call should be avoided when performing RAII, because it will have the Database keep a smart pointer to the Transaction that will prolong the call to abort() if an exception is thrown.

Database Environments

Database environments are convenient for configuring a group of Berkeley databases that will be used together. Using database environments together with the Concurrent Data Store mode with multiple

Listing 11. STLdb4 and Database Environments

```
#include <iostream>
#include <STLdb4/stldb4.hh>

using namespace STLdb4;
using namespace std;

int main( int argc, char** argv )
{
    string dbenvpath = argv[1];
    fh_env dbenv = new Environment( dbenvpath );
    dbenv->setDefaultOpenFlags(
        DB_CREATE | DB_INIT_CDB | DB_INIT_MPOOL );
    Environment::setDefault( dbenv );

    fh_database db = new Database(
        DB_BTREE, dbenvpath + "/foo.db" );
    db["bar"] = argv[2];

    fh_database db2 = new Database( dbenv );
    db2->create( DB_BTREE, dbenvpath + "/foo2.db" );
    db2["key"] = (string)"value_" + argv[2];

    return 0;
}
```

database files allows you to have multiple applications all reading and writing to the database files, and Berkeley DB takes care of locking to make sure that the files don't become corrupt.

The default database environment in STLdb4 is effectively a null environment. New database environments are created using the Environment class. The static Environment::setDefault() method can be used in applications using a single database environment to avoid having to pass the database environment object to the Database constructor.

The code shown in Listing 11 uses a database environment to protect two database files from simultaneous update by multiple running processes. First, a new database environment is created and set to use the Concurrent Data Store mode. This database environment is set to be the default STLdb4 environment. The first Database object is created using the default database environment; the second Database object is created by specifying the database environment explicitly and opening the database file as a separate call.

Other Things of Interest

The ordering of elements in the database can be changed with Database::set_bt_compare() using either a function pointer or a Loki functor object. For details on Loki functors, refer to the *Modern C++ Design* book (see the on-line Resources). As the

Listing 12. STLdb4 and Database Environments

```
fh_database db = new Database();
Database::m_bt_compare_func_t tmpf
    = getInt32Compare();
db->set_bt_compare( makeReverseCompare( tmpf ) );
db->create( DB_BTREE, "/tmp/play.db" );
```

comparison function is a relatively low-level operation, no implicit conversions happen for this, and you must compare two Dbt values. A collection of functions for numeric comparison, such as getInt32Compare() and string comparison with and without case sensitivity using getCISCompare(), are available in STLdb4. The ordering of a comparison functor can be reversed by passing it to makeReverseCompare() to create a new functor. These operations must be performed before the database is open, so you have to use the open() or create() calls and the non-opening Database constructor as shown in Listing 12.

Increasing the default Berkeley DB cache size using Database::set_cachesize() can improve read-only database performance significantly.

Future Directions

A template subclass of Database taking parameters similar to std::map<> would be nice. There would need to be some extra work to allow both the key and value to be (de)serialized on demand perhaps by assuming that both can be Boost serialized. ■

Resources for this article: www.linuxjournal.com/article/9512.

Ben Martin has been working on filesystems for more than ten years. He is currently working toward a PhD at the University of Wollongong, Australia, combining Semantic Filesystems with Formal Concept Analysis to improve human-filesystem interaction.

axigen **ISP EDITION**

The ISP Solution of Choice for a Carrier Class Mail Server

CLUSTERING SYSTEM High Availability (RHCS integration) Distributed mailboxes (LDAP routing)	ADVANCED SECURITY Message Acceptance/Routing policies Antivirus & Antispam integration Anti-impersonation function	REPORTING ENGINE (over 100 personalized reports) ONLINE BACKUP AND RESTORE SYSTEM
--	--	---

www.axigen.com/Lj
 Visit now to start your free trial!

Creating SELinux Policies Simplified

SELinux does not have to be as hard to manage as many people may think.

IRFAN HABIB

Modern Linux distributions ship with a plethora of security features and tools, and one of the most important features added to the kernel has been the inclusion of SELinux.

SELinux solves one of the most challenging problems in security. How to control trusted users or processes. Trusted users, such as root in the *nix domain, have unlimited and unrestricted access to a system. They should have, as this account is supposed to be used only by the system administrator; however, this leads to a problem. What happens if the root account itself gets hacked, leaving the hacker with complete control of the system. Malicious users are not the only problem; mis-configuration of a security tool, such as iptables, can have a profound effect as well. Besides this, imagine a security vulnerability is discovered for a service you deployed on your server, and no patch is made available in a timely manner. In this case, your system is vulnerable. SELinux secures Linux systems from these sorts of security issues by implementing mandatory access controls (MACs) in the Linux kernel. SELinux is based on the Flask security architecture. Discussion of the Flask architecture is avoided in this article, as excellent documentation about it is easily available on the Internet.

To understand mandatory access control, we must go through the currently deployed security model called discretionary access control (DAC). In a DAC system, access to objects is restricted based on their classifications. This type of control is discretionary in the sense that a subject with a certain set of access permissions is capable of passing those permissions on to another subject. For example, any program you run while logged on as a certain user has the same access rights that you have. Rights are set by another user (for example, root).

Any particular permission (read, write, execute and so on) can be thought of as a two-dimensional graph with users on one axis and objects on another. In essence, DAC systems check the validity of credentials presented to them against stored information.

As mentioned, the SELinux security model is mandatory access control, or MAC. This controls access in a different manner. Whereas DAC security models are authentication-based, MAC systems rely on authorization, not only of users but also of each object loaded by the system.

A MAC system controls objects individually and makes decisions on the rights and/or permissions of objects based on a security policy, which can define what rights the object should be accorded, based on different variables.

An example of how discretionary versus mandatory access control styles could affect the operation of a computer is a Python script. If the script allows an external entity to insert and execute malicious code on a computer system under a DAC system, the malicious code now has the same access rights as the code that executed it—the Python script.

A MAC system can restrict the rights of a certain process to only the resources needed for normal operation. A Python script may create a process (or it may be forbidden), but that process might not have the same set of permissions as the process that created it. Thus, the MAC approach is seen as secure.

A Brief Look at SELinux Internals

In SELinux, the security policy configuration is defined in a text file written in the m4 language. It is compiled when the security policy is finalized and, at boot time, loaded into the memory. Only the security server can make policy decisions on the permissions of an object.

Security policy enforcement is done by components called object managers, which receive requests from client objects, submit queries to the security server and enforce the resulting decisions.

The SELinux implementation of the security server uses a combination of two security paradigms, called Type Enforcement (TE) and Role-Based Access Control (RBAC).

Type Enforcement

Type Enforcement makes security decisions based on what kind of object requests the permissions. For example, object types could include a regular file, a directory, a process or a socket. Type Enforcement is an object labeling system that, combined with access mapping (from the domain of the object requesting permission and to the type of the object requested), returns a decision that defines the permissible actions of the object.

Role-Based Access Control

Role-based access control assigns permissions to objects in a computer system based on the role they play within that system. In practice, this means a process would have its permissions based on its parent process, the user logged on at the time and any number of other variables.

How processes, filesystem objects and sockets communicate with each other is defined by the security policy. In particular, the security policy governs how different types and roles may interact, along with any specific rules.

At present, SELinux provides binary compatibility with existing applications and source compatibility with kernel modules. The current implementation of SELinux is x86-specific.

Getting and Installing SELinux

SELinux is included in many distributions nowadays, and even if it does not come with the main distribution, distribution-specific packages usually are available for popular distributions. The sources for SELinux are available from www.nsa.gov/selinux/code/download0.cfm.

Installation of seedit

seedit is a user-friendly Webmin-based tool that enables an administrator to administer SELinux policies from a Web browser. The seedit Webmin interface allows users to perform every action they can normally perform by writing manual scripts in m4 in a point-and-click environment. seedit is available from seedit.sourceforge.net.

Once installed, the SELinux policy is located in `$(SELINUX)/(seedit-something)/policy/policy.conf` (where `$(SELINUX)` is the root directory of your SELinux installation—usually it is `/etc/selinux/`).

Let's familiarize ourselves with the seedit interface. Fire up a browser,

Why ADD Value, When You Can Multiply It?

ACMA is a leading provider of custom manufacturing solutions. With world-class engineering, ISO 9001 certified manufacturing, nationwide logistics, and 24/7 technical support, ACMA will not just ADD value to your business.

ACMA will multiply your business success!

Call now for a personal consultation
1.800.786-6888 or visit www.acma.com

point it to <http://localhost:10000> and go to the System→SELinux Configuration section.

You will see six icons labeled configure ACLs, define domain transitions, define relationship between users and roles, create new domain/roles, delete domains/roles and update configuration.

Configure ACLs

Here you can define virtually all access control to nearly every object in the system. This includes allowing/disallowing read, write or executable files to entire directories or individual files and allowing/disallowing access to networking capabilities. In this section, you also can define ports in which the specific application can work. Let's say we assign port 80 to Apache; if it is started to work on port 81, SELinux terminates this process.

IPC access control can be defined in this section also. You can define what kind of IPC mechanism this particular application can use and with which applications the particular application can communicate.

Other access controls that can be defined in this section include various administrative access controls, such as kernel communication privileges, SELinux operations, process information retrieval and so forth.

Defining Domain Transitions

Should it be Domain Transitions or Transactions throughout? In the domain transition section, you can define which processes can spawn the current application—for example, by default, the seedit policy defines this domain transition for MySQL: kernel→init→mysqld. This means the kernel can start init, and init in turn can launch the MySQL daemon. If the application has a daemon associated with it, domain transitions should be defined or the daemon will never be able to start.

So in this section, the user can define domain transition, alter existing ones or remove them altogether.

Define Relationship between Roles and Users

Roles are privileges that an object (like users) can have on a system. For example, there may be a role that allows access to all files in the system. Thus, in this section, users can associate roles with specific users on the system.

The next two sections are self-explanatory. They allow you to define a new domain/role and remove a domain/role. Defining a domain/role is the first step when adding access control rules for a new application or user.

The last section is the update configuration section; it allows the user to update the policy and recompile it and load it.

seedit comes with a plethora of predefined policies, which cover every popular server/daemon in a Linux system—from the kernel to the MySQL daemon.

Defining a New Policy for a Daemon Application

Defining an SELinux policy for a daemon is an iterative process. The first step is to "register" the daemon with SELinux, by declaring a domain/role in the seedit Webmin interface. This is done in the Add Domain/role section discussed previously. The convention is that domains should end with a `_t`, and role definitions should end with an `_r`. So, there could be a role, such as `admin_r`, and a domain, such as `mysqld_t`.

Defining domain transitions is another important step. Here you have to allow the parent processes of the daemon to spawn the concerned daemon. Usually if the daemon activates at startup, you need to define a domain transition from init to the concerned daemon.

Define the initial access control list for the daemon. Now, it is not possible for the user to define an exhaustive ACL for the daemon when it is initially installed, rather what is usually done is that an initial ACL list is defined, which is very restrictive. Every time the application tries to access an object and is not allowed to do so, an access violation message occurs; such messages can be accessed from `/var/log/messages`. As you use the daemon, SELinux will report some violations. Considering that the daemon for which we are trying to write rules is called foobar, the violation mes-



Service

Quality

Customization

Acma[®]
We Manufacture For You!

See Us at West 2007 San Diego January 31 - February 2, 2007 Booth#1802

Toll Free: 800.786.6888 **WWW.ACMA.COM**

© 2006 Acma Computers Inc, 3393 West Warren Avenue, Fremont, CA 94538. All rights reserved.

sages will look like this:

```
javc: denied { write } for pid=7279 exe=/usr/bin/foobar comm=ifup
name=dhclient-eth0.conf dev=hda12 ino=57400
scontext=system_u:system_r:foobar_t tcontext=system_u:object_r:etc_t
tclass=file
```

This violation states that a write attempt was made to the file `dhclient-eth0.conf` by our application. This violation can be removed by giving the application write access to the `dhclient-eth0.conf` file. This violation can be addressed by going to the Configure ACL section of the `seedit` Webmin interface, under the File ACL section, and then browsing to the place where this file exists, and giving the application `foobar` write access to it.

Another violation might look like this:

```
javc: denied { create } for pid=7279 exe=/usr/bin/foobar
scontext=root:system_r:foobar_t tcontext=root:system_r:foobar_t\
tclass=udp_socket
```

This violation reports that the application tried to create a UDP socket and was denied. To remove this violation, we can simply add networking features to the access control of the `foobar_t` domain.

This can be addressed by going to the Configure ACL section and in the Network ACL section, and selecting Allow Network for the domain `foobar_t`.

All access violations can be addressed in the Configure ACL section of the `seedit` Webmin interface.

After every policy update, reload it via the Update Policy option in the `seedit` Webmin interface and restart the `dæmon`. As other violations occur, update it again and so on, until you get nearly no violation messages.

There are other ways to generate access rules for an application with the help of `audit2allow` tool. But using it can lead to general rules, which can lead to security problems. You always can refine the rules the tool generates, however.

When creating or updating policies, make sure you have set your SELinux installation to permissive mode. SELinux has three modes: enforcing, disabled and permissive. In enforcing mode, all access controls are enforced according to the defined policy. In permissive mode, the policy is not enforced; however, violation messages are shown when any violation to the policy occurs. Disabled mode completely disables SELinux. ■

Irfan Habib is an undergraduate student of software engineering at the National University of Sciences and Technology, Pakistan. He has been deeply interested in Free and Open Source software for years.

Is your Linux/Unix-based server passing on viruses?



Our CSAV Interceptor will allow you to protect your Windows users from malware.

Scan email on Linux/Unix-based servers with Authentium's CSAV Interceptor for Sendmail.

If users have to wait for a warning, it could be too late.

Need we say more?

Solutions for Exchange, Clearswift and Postfix also available.

www.authentium.com/linuxjournal (800) 423-9147 interceptor@authentium.com

Integrating PHP and Perl

PHP and Perl are both so powerful, they can even run each other.

IRFAN HABIB

Perl is a language often associated with text processing and CGI. PHP is a language often associated with dynamic Web pages. Both are very popular with Web developers. Often, each of these languages is used at the expense of the other. Hard-core Perl developers would love to develop everything in Perl, and PHP developers tend to stick with PHP.

As usual in the Open Source world, there is a lot of zealotry between users of each language. If you think that one of these languages is perfect and the other is lame, this article is not for you! This article is for those who take a more pragmatic approach and use what works best for them. Each language has its strengths and limitations. Personally, I use both languages at work and at home. With time, I have discovered which language is best for which tasks and try to integrate the strengths of each language as much as possible to complete my work quickly.

Perl is extremely good at system administration and extensive data processing, among other things. This means, if you want to do some extensive processing on a text report, Perl would be preferable, as it provides handy regular-expression-enabled text comparisons, which make it so much easier to search through a report. Perl also has extensive string manipulation features. Perl, by virtue of being older than PHP and having an extensive community, has thousands of extensions archived in CPAN, which allow one to do virtually anything with the language, conveniently. From XML processing to writing to parallel port devices, CPAN includes everything. CPAN is the reason Perl continues to be useful to a large number of developers to date. Although it is not impossible to do everything described here with PHP and a mixture of other languages, it's simply more convenient with Perl.

PHP is extremely good at integration with Web pages and databases. PHP integrates nicely with static HTML Web pages. That's why it's so popular and has had more visibility than Perl in recent years. It has mature support for numerous popular free or non-free databases and supports MS SQL (MSSQL) Server better than any other open-source language. From personal experience, I have tried at least two CPAN extensions for Perl to get it to work with an MSSQL installation, but with limited success. However, PHP has seamless support for MSSQL and uses it as natively as MySQL.

I was recently involved in a project in which nearly the entire project was in Perl. However, a tiny bit of code needed access to an MSSQL server. I knew how simple it was in PHP to work with MSSQL, and I did not want to go through the pain of setting up my Perl installation for MSSQL. That's why I searched the Internet for a way to integrate both languages in a manner that would allow me to use the best parts of each language and produce a coherent solution. And, I found the PHP::Interpreter CPAN module. PHP::Interpreter was perfect. It enables the complete integration of the two languages to an extent that one starts to believe that both are mere extensions of each other. PHP::Interpreter, as this article shows, allows you to use PHP's mature support for databases and other features natively in Perl, and also to use Perl's vast number of CPAN modules to extend your PHP programs.

According to AnnoCPAN, the module's main function is to encapsulate an embedded PHP5 interpreter. It provides proxy methods (via AUTOLOAD) to all the functions declared in the PHP interpreter, transparent conversion of Perl data types to PHP (and vice versa), and the ability for PHP to call Perl subroutines similarly and access the Perl

symbol table. The goal of this package is to construct a transparent bridge for running PHP code and Perl code side by side.

To demonstrate the power of this module, we code two examples to show each side of the PHP::Interpreter, integrating Perl with PHP and integrating PHP with Perl. Each example shows areas in which both languages complement each other nicely to produce powerful code.

EXAMPLE 1 Integrating PHP with Perl

In the first example, we create an application to monitor failed logins through SSH to our system. SSH often is targeted by script kiddies and malicious users to compromise a system and gain access to it. The script identifies the IPs of the offenders, blocks all incoming packets from using iptables and, finally, logs them in to an MS SQL server database. We use Perl to do what it's best at—processing log files. It will continuously monitor the `/var/log/messages` file, which the SSH daemon uses to log failed login attempts. To monitor a log file continuously, we use the CPAN extension `File::Tail`. To support writing to MS SQL Server transparently, we implement this portion in PHP and show how the two languages can be integrated seamlessly and used in scenarios where both complement each other.

Setting Up PHP::Interpreter

Setting up PHP::Interpreter is basically a standard Perl module installation procedure. You can get it from search.cpan.org/dist/PHP-Interpreter. Unpack it, and create the Makefile:

```
perl Makefile.PL
```

Compile it:

```
make
```

And, install it:

```
make install
```

You can do an additional:

```
pod2html interpreter.pm > interpreter.html
```

and keep the documentation file for future reference.

We also use the CPAN module `File::Tail`, which allows us to monitor a log file continuously. You can get this module from search.cpan.org/dist/File-Tail.

Unpack it, and create the Makefile:

```
perl Makefile.PL
```

```
make
```

```
make install
```

Now, fire up a text editor, and start coding:

```
1. use PHP::Interpreter;
2. use File::Tail;
```

```

3. use threads ('yield', 'stack_size' =>64 * 4096, 'exit'
=>'threads_only');
4. use Thread;
5. my $php = PHP::Interpreter->new;
6. my $ref=tie *FH,"File::Tail",(name=>'/var/log/messages');
7. while (<FH>)
8. {
9. if($_~/sshd/) #checks for message from sshd
10. {
11. if($_!/Failed password for/) #check for a failed password attempt
12. {
13. $ind = rindex($str,'from');
14. $rind = rindex($str,'port');
15. $ip = substr($str,$ind+4,$rind-$ind-4);
16. $thr = new Thread \&writems, $ip;
17. $thr->join();
18. }
19. }
20. }
21. sub writems
22. {
23. `iptables -I INPUT -s $ip -j DROP`
24. $php->include("**writems.php**");
25. $php->writeIP('$sqlserver', 'sshwatch', 'sshusr', 'sshpass', $_[0]);
26. print $php->eval("echo Succeeded!");
27. }

```

In a separate file, write the following script (the file should be named `writems.php`):

```

1. <?php
2. function writeIP($dbhost,$dbname,$dbuser,$dbpass,$ip)
3. {
4. $conn = mssql_connect($dbhost,$dbuser,$dbpass)
5. or die("Couldn't connect to SQL Server on $dbhost");
6. $db = mssql_select_db($dbname, $s)
7. or die("Couldn't open database $myDB");
8. set_time_limit(0);
9. $squery="insert into sshwatch(currentdate,ip)
10. values('".date('Y/m/d')."','".$ip."')"; mssql_query($squery);
11. }
12. ?>

```

To run the application, simply run the Perl script:

Perl scriptname

In Line 25, you need to fill in the correct settings for your MSSQL server installation. You also need to have a PHP installation with support for MSSQL. This is usually done by passing the switch `-with-mssql` during the compilation of PHP. Some distributions also require you to install FreeTDS, which is used by PHP to access MSSQL.

Now, let's review some specific portions of the code. To use the `PHP::Interpreter` in your code, declare its use, as in line 1. To create a new instance of PHP interpreter, do as is shown in line 5:

```
my $php = PHP::Interpreter->new;
```

As with object-oriented programming, you now can invoke methods on the `$php` object to achieve interoperability with PHP. The above code shows two functions provided by the `PHP::Interpreter` for interoperability. In line 24, we are calling the `include()` function, which includes a PHP script file to the environment, and you can call functions defined in the file

natively from the object. We do the same with `writelP` in line 25, which is a PHP function declared in `writems.php` on line 2 of the `writems.php` listing. The `Eval` function of the `$php` object allows you to execute a specific PHP instruction, as with a live interpreter. The instruction is interpreted, and the return may be stored into a variable or used directly, as in line 26. As you can see in the above program, `PHP::Interpreter` provides an object-oriented mechanism for completely integrating the two languages. This integration is achieved with only two lines of code: the initial use statement and the instantiation of the object. `PHP::Interpreter` is not only about calling functions and procedural programming, it also works with object-oriented PHP. This is how you can instantiate an object of class defined in a PHP:

```
my $instance = $PHP->instantiate('PHPclass', @args);
```

The instance is stored in `$instance`, and any arguments are passed to the class' constructor.

EXAMPLE 2 Integrating Perl with PHP

The biggest advantage of Perl/PHP integration is PHP's ability to access Perl CPAN modules. There are CPAN modules for almost everything that can be done via software; you can use `PHP::Interpreter` in PHP to call CPAN modules to extend a PHP application to do anything, which is not native to PHP—for example, it enables you to write to IO ports. Writing to IO ports has been the exclusive domain of C/C++ programs, but with `PHP::Interpreter`, even a mere scripting language can have the capability to write to IO ports. The example that follows shows how to use Perl code with PHP, but first, we discuss the features of `PHP::Interpreter` that allow PHP/Perl integration.

The PHP interpreter, invoked via `PHP::Interpreter`, has a special class that allows PHP to Perl communication. Create an instance of the class via this call in PHP:

```

1. <?php
2. $perl = Perl::getInstance();
3. ?>

```

The new `$perl` object allows you to evaluate specific Perl instructions in PHP, such as:

```

1. <?php
2. $perl = Perl::getInstance();
3. $perl->eval(q^
4. print "Executing Perl code in PHP\n";
5. ^);
6. ?>

```

Similar to Example 1, where we called a PHP function in Perl, you can call Perl subroutines in PHP. All subroutines defined in the Perl program, which instantiated the `PHP::Interpreter` instance, can be invoked like this (I will provide a more detailed example shortly):

```

1. <?php
2. $perl = Perl::getInstance();
3. $return = $perl->call('sub', @args);
4. ?>

```

And, of course, you can get and set variables from the Perl file that instantiated the `PHP::Interpreter`; however, only package variables, not lexical variables, are supported.

Let's look at a practical application of PHP/Perl integration—for example, a snippet of Perl code that uses the Babel Fish CPAN module. (Babel Fish is a piece of software that allows you to translate text between different languages. To learn more about Babel Fish, go to <http://babel.altavista.com>.)

The PHP program calls the translate function, which will be implemented in Perl, to translate a string in English to German and retrieve the output.

To install the Babel Fish CPAN, go to search.cpan.org/CPAN/authors/id/D/DM/DMUEY/AltaVista-BabelFish-v42.0.1.tar.gz, and install it with the standard installation procedure, as shown previously in this article.

AltaVista::BabelFish also has some prerequisites, such as Class::Std and Class::Std::Util. These need to be downloaded and installed for Babel Fish to work:

```
1. use AltaVista::BabelFish;
2. use PHP::Interpreter;
3. my $p = PHP::Interpreter->new();
4. $p->include("phpscript.php");
5. my $val = $p->invoke();

6. sub translate
7. {
8. my $phish = AltaVista::BabelFish->new({ source => $_[0], target =>
   $_[1] });
9. return $phish->translate($_[2]) or die $phish->get_errstr();
10. }
```

The phpscript.php file contains the following:

```
1. <?php
2. function invoke()
3. {
4. $perl = Perl::getInstance();
5. $string = $perl->call('translate', 'en','de','Translate this for me');
6. print "Translated string: $string\n";
7. }
8. ?>
```

Let's look at this piece of code in more detail. In line 4 of the PHP program, we are creating an instance of the Perl class using Perl::getInstance(). This is the special class inserted by the PHP::Interpreter dynamically into the environment to achieve PHP to Perl integration.

In line 5, we then use the class object, \$perl, to invoke a function called translate, which is defined in line 6 of the Perl program, and we pass the arguments accordingly. The subroutine translate is invoked from the Perl script, and the translation is done via the Babel Fish module. The translated string is returned to PHP and printed via the print statement. Although this is a rudimentary example, the entire script can be extended to provide runtime translation for viewers of a dynamic Web page generated from PHP. With CPAN and the PHP::Interpreter, the possibilities of what can be achieved in PHP are bounded only by the developer's imagination.

You can use the PHP Perl class for object-oriented Perl as well. Invoke a Perl object via the new() function, as follows:

```
1. <?php
2. $perl = Perl::getInstance();
3. $instance = $perl->new('perlclass', @args);
4. ?>
```

The first argument to the new() method, in line 3, is the name of the class, and additional arguments are passed to the constructor of the class.

Conclusion

This article shows both sides of the PHP::Interpreter: using PHP in Perl and Perl in PHP. The module essentially allows a PHP programmer to extend the capabilities of PHP to enable it to do anything that CPAN allows Perl to do. It also allows a Perl programmer to use those features in PHP that are not yet mature or not implemented in Perl. By no means have I covered all of the PHP::Interpreter, and readers are encouraged to explore the official CPAN documentation of PHP::Interpreter. ■

Irfan Habib is an undergraduate student of software engineering at the National University of Sciences Technology. He has been deeply interested in Free and Open Source Software for years. He often comes across tasks for which he needs to pull together a solution really quickly, and Perl and PHP usually allow him to do that. He can be reached at irfan.habib@niit.edu.pk.

NEW!

Tiny WiFi Controller
boots Linux in 1.1 seconds

\$129
CPU board only


\$249
as shown




quantity
discounts
start at
10 units

 **200 MHz CPU**

- TS-7400 CPU board
- Low power, low heat, long life
- Up to 128MB on-board Flash
- Up to 128M SDRAM
- SD Flash Card socket
- 1 external USB port
- 1 10/100 Ethernet
- 802.11g internal WiFi option
- One piece, rugged aluminum enclosure option measures 1.1" x 4.9" x 3.1"

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

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

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



We use our stuff.

visit our TS-7200 powered website at

www.embeddedARM.com/7400wifi

(480) 837-5200

Painless Thumbdrive Backups

Exploit udev rules to back up your Flash drive daily or every time you insert it.

ANDREW FABBRO

Raise your hand if you've ever lost (or worried you'd lost) a USB thumbdrive. You spent hours fruitlessly searching the house, and then as you opened the washing machine door, it suddenly dawned on you that perhaps you didn't check your pockets thoroughly when you put this load in.

Fortunately, you have a backup of all the data, right? You religiously mount the drive and copy the data to a backup directory on a regular schedule, no?

That sounds an awful lot like drudgery to me too, and I got into computers to avoid boring work. Naturally, it's a lot more fun to spend some time working out the perfect method for painless thumbdrive backups.

What do I mean by painless? How about a system where you can walk up to your Linux box, plug in the drive, wait for a "backup complete" sound, unplug and walk away? Perhaps a system that keeps its backups orderly (say, the last seven copies)? Oh, and it should handle encrypted thumbdrives as well. And, if you need to recover, it should do both whole-volume replacement and per-file restores.

Not a problem. The key to this system is using udev rules and a simple shell script. The tools already are on your system. In this example, I use a CentOS 4.3 system, though any Linux distribution with a 2.6 kernel should work.

udev to the Rescue

udev is the modern device manager for Linux, replacing the 2.4 kernel's devfs. udev handles all device mapping, including hot plugging of devices. One of its coolest features is it lets you write your own event rules. This article shows you how to craft a rule that automatically fires when you plug your USB thumbdrive in to the system.

These rules are stored in `/etc/udev/rules.d` (if you're using a different Linux distribution, check `/etc/udev/udev.conf` for the `udev_rules=` line, which should point to the rules directory). You can place whatever udev rules you want as text files in this directory, and udev picks them up immediately for use without requiring a reboot.

udev

This article has scratched only the surface of what you can do with udev rules. Any type of hot-plug event can fire a rule that can do almost anything. For example, you can write rules to mount devices automatically, copy pictures off a digital camera or set up a network link. udev's rules language provides great flexibility, including printf-like wild cards and the ability to set permissions.

The best overview for writing your own udev rules is Daniel Drake's "Writing udev Rules", which can be found at www.reactivated.net/writing_udev_rules.html.

How to Identify Your Device

To write a udev event rule, you first need a unique way to identify the USB device. Most thumbdrives have serial numbers, though not all. Fortunately, even with thumbdrives that do not have a serial number, you can craft udev rules for them.

I use two thumbdrives as examples: a JetFlash JF110, encrypted with TrueCrypt, and a Corsair Flash Voyager. The JetFlash has a serial number; the Corsair does not.

Plug your thumbdrive in, and `cat /proc/scsi/usb-storage/*`. You should find an entry for it similar to this:

```
Host scsi5: usb-storage
  Vendor: Unknown
  Product: USB Mass Storage Device
Serial Number: 85a5b1f2c96492
  Protocol: Transparent SCSI
  Transport: Bulk
  Quirks:
```

If you have a serial number, skip forward to the "Writing the Rule" section of this article. If you see "None" for the Serial Number, you still can identify the device by using `udevinfo`. Follow these steps:

1) Look at `dmesg`'s output. Typical output is as follows:

```
usb-storage: waiting for device to settle before scanning
  Vendor: Corsair   Model: Flash Voyager   Rev: 1.00
  Type: Direct-Access
  ANSI SCSI SCSI device sde: 2031616 512-byte hdwr sectors (1040 MB)
[...]
sde: assuming drive cache: write through
sde: sde1
Attached scsi removable disk sde at scsi12, channel 0, id 0, lun 0
Attached scsi generic sg4 at scsi12, channel 0, id 0, lun 0, type 0
```

This tells you that `/dev/sde` is the device assigned.

2) Now, run:

```
udevinfo -a -p $(udevinfo -q path -n /dev/sde)
```

and examine the output. Look for these lines:

```
BUS=="scsi"
SYSFS{model}=="Flash Voyager  "
SYSFS{vendor}=="Corsair  "
```

Writing the Rule

Now, with either the serial number or the vendor/model combo, you can write the rule. The rule creates a symlink for the device in the `/dev` tree, for example, `/dev/corsair_drive`, and then calls the script `/usr/local/bin/backup-thumb.sh`, which I'll get to in a moment.

Become root (`su -`), and create a text file in `/etc/udev/rules.d` called `95.backup.rules`. You can use a number other than 95, but keep in

mind that udev processes rules in alphanumeric order, and it's better to have local rules processed last.

If you have a serial number, type a rule like this (all on one line) into the file, and save it:

```
BUS="usb", SYSFS{serial}="85a5b1f2c96492", SYMLINK="jet_drive",  
RUN+="/usr/local/bin/backup-thumb.sh jet_drive "
```

If you're using vendor/model identification, your rule would look like this:

```
BUS="scsi", SYSFS{vendor}=="Corsair ", SYSFS{model}=="Flash Voyager ",  
SYMLINK="corsair_drive", RUN+="/usr/local/bin/backup-thumb.sh  
corsair_drive"
```

Note that you can string as many SYSFS{} entries together as you need to identify the drive uniquely. Your rule now fires every time you plug in your thumbdrive.

Note: if you have other rules for a device, udev executes the rules in sequence from top to bottom.

Set Up the Backup Script

backup-thumb.sh is the engine that backs up your thumbdrive. Our rule calls it, giving the name of the device (the SYMLINK) as its only argument. Everything else is configured in the CONFIG section. The backup script is shown in Listing 1.

Put this script in /usr/local/bin/backup-thumb.sh, and remember to chmod +x it. Next, edit the CONFIG section—the parameters are as follows:

- BACKUP_DIR: where you want the backups to go.
- GENERATIONS: how many days of backups to keep. Backups will be numbered 0 (most recent) to the limit you enter (oldest). Keep in mind that you need to have enough storage space for this many backups. If you are backing up a 1GB fob and set GENERATIONS to 7, backups will consume 7GB of space.
- BACKUP_ONCE_DAY: if you plug and unplug your fob multiple times a day, you probably won't want to back it up each time. backup-thumb.sh uses a tag file so that it backs up only once per day. If you want to change this so it runs a backup every time you plug in a thumbdrive, set BACKUP_ONCE_DAY to 0.
- SOUND: in this example, I've chosen a sound from the KDE distribution, but any WAV file will work. You easily can modify the script to use madplay instead of aplay and use an MP3 file as your completion sound.

How It Works


backup-thumb.sh sleeps for ten seconds on startup, because it must wait for the kernel to finish scanning the thumbdrive. If you plug in a thumbdrive and type dmesg, you'll see a "waiting for device to settle" message while this happens. Ten seconds for the kernel scan should be sufficient even for older machines.

Next, backup-thumb.sh sets permissions tightly so that only root can read the backups. Otherwise, some nefarious person could copy your backup to a different machine and mount it there.

The script executes a simple dd (bit-for-bit copy) of your thumbdrive to a backup file. This works whether the device is encrypted or not. When it's finished, it plays a noise you will hear on your computer's speakers. On a USB 2.0 port, backing up a 1GB thumbdrive takes about one minute.

Linux[®] the Negus Way.

From Fedora[™] to Knoppix, these are the
latest resources you'll need to succeed.



Linux Bible 2007 Edition 0-470-08279-8

Fedora 6 Bible 0-470-08278-X

Buy your copies today at
www.wiley.com/go/negus



be **FREE**



Join us at member.fsf.org



Listing 1. Backup Script

```
#!/bin/bash

# Thumbdrive backup script from Linux Journal

# #####
# CONFIG section

# where you want the backups to be kept
BACKUP_DIR=/backups/thumb

# how many old backups to keep
GENERATIONS=7

# backup only once a day
# set to 0 if you want a backup run every time
# you insert your thumbdrive
BACKUP_ONCE_DAY=1

# completion sound to play when backup is done
SOUND=/usr/share/sounds/KDE_Beep_ClockChime.wav

# END CONFIG
# #####

# main program

# wait for device to settle
sleep 10

# make sure no one will be able to copy our backups
umask 077

# check the directory
DEVICE=$1
if [ ! -d ${BACKUP_DIR} ] ; then
    mkdir -p ${BACKUP_DIR}
fi

# only backup once per day
if [ ${BACKUP_ONCE_DAY} -gt 0 ] ; then
    DIDTODAY=${BACKUP_DIR}/${DEVICE}.did_today
    find ${BACKUP_DIR} -name ${DEVICE}.did_today -a -mtime +1 -delete
    if [ -f ${DIDTODAY} ] ; then
        exit
    else
        touch ${DIDTODAY}
    fi
fi

# rotate backups
cd ${BACKUP_DIR}
let GENERATIONS=${GENERATIONS}-1
while [ ${GENERATIONS} -ge 0 ] ; do
    let NEWFILE=${GENERATIONS}+1
    if [ -f ${DEVICE}.backup.${GENERATIONS} ] ; then
        mv -f ${DEVICE}.backup.${GENERATIONS}
        ${DEVICE}.backup.${NEWFILE}
    fi
    let GENERATIONS=${GENERATIONS}-1
done

# do the backup
dd if=/dev/${DEVICE} of=${BACKUP_DIR}/${DEVICE}.backup.0 > /dev/null 2>&1

# notify that we're done
aplay ${SOUND} > /dev/null 2>&1
```

How to Recover

If you lose your thumbdrive and want to restore your backup to its replacement, simply dd the backup image to the new thumbdrive, like so:

```
dd if=corsair_drive.backup.0 of=/dev/corsair_drive
```

Or, if you want to grab only some files from the backup, do the following:

```
mkdir /mnt/thumb
mount -o loop corsair_drive.backup.0 /mnt/thumb
```

You now can copy the files from /mnt/thumb.

If you're using TrueCrypt to encrypt your thumbdrive, you can mount the backup image in much the same way:

```
truecrypt corsair_drive.backup.0 /mnt/thumb/
```

That's about as painless as we can make thumbdrive backups. If you're too lazy to plug your drive in and come back when it beeps...well, stay away from laundromats! ■

Andrew Fabbro has become an Oracle DBA; however, he still has root at home and he welcomes your comments sent to andrew@fabbro.org.

STOP Crackers Dead!

ZyWALL Security Appliances Protection for your Business



NOWTHOR
CORPORATION **1-877-371-5778**
shopping.nowthor.com

Long Live the Freedom of Linux

A window of opportunity awaits for Linux that is begging for proper freedoms.



Nick Petreley, Editor in Chief

Microsoft is running scared, and has many good reasons to be afraid. Vista is destined to be a disappointment, an expensive one for Microsoft, OEMs and customers. I'm not saying it will flop. Vista will find its way onto many desktops because it will be preloaded on computers. But aside from one single Microsoft groupie, I don't know anyone who cares about Vista, let alone anyone who is excited about its release. Microsoft knows how much the Windows line has been botched. That's one reason why Microsoft is now trying to make money off every copy of SUSE sold by Novell. For more details, see my Web article on the *Linux Journal* Web site, "A Five Year Deal with Microsoft to Dump Novell/SUSE" (www.linuxjournal.com/node/1000121). Praise God, Red Hat flatly refused to sign a similar deal with Microsoft.

This inevitable botch job gives Linux an invaluable window of opportunity to grab desktops away from Windows. Why migrate to an expensive, bloated, poorly designed Windows when it's just as easy (or hard) to migrate to Linux on the desktop?

The best way to seize this opportunity is to make Linux distributions even more free, as in freedom, than they are now. Here are

the freedoms I propose:

1) As stated in the Web article referenced above, make all your computers free of any Novell/SUSE software.

Don't pay Novell to funnel protection money to Microsoft so Microsoft won't sue you. Make Microsoft sue someone over patent infringements. I say "someone" because that someone really isn't likely to be customers. Such a bad PR move would bury Microsoft, and everyone at Microsoft knows it. Microsoft might sue *someone* over some patent or another, but let's get these issues in and out of court once and for all and have the issues settled for eternity. You can't make that happen if you pay Novell to support Microsoft's protection racket.

2) Make Linux free of anything that Microsoft can claim is a patent infringement. Meanwhile, scour the code for anything that Microsoft might claim infringes upon its patents and replace it. OpenGL is a good place to start. Microsoft bought the patent for OpenGL from SGI in one of its many bail-out deals.

3) Make Linux free with every computer. Linux distributors, both commercial and free, need to band together and pressure hardware OEMs to preload a preconfigured version of Linux. This is a difficult step because it calls for selfless cooperation between distributors, knowing that a rising tide lifts all boats. Red Hat should be content if its part of the joint effort causes Dell to preload Ubuntu instead of Red Hat or Fedora. Once Linux becomes a standard preloaded operating system, that is the time for individual distros to start fighting to sway the OEMs toward their respective products. Until then, they must fight as a group to get anything preloaded on the systems.

4) Make the Linux desktop and common applications geek-free.

I'll be optimistic and assume more Linux will get preloaded. In such a case, far more computer-illiterate people will use Linux. So rip the geek out of the desktop and the handful of the most popular end-user-oriented applications. Take these applications and

bring all the most commonly used features to the surface. The photo management program LPhoto is an excellent example of one way to do this right. Things like red-eye reduction, print photo, e-mail photo and other common functions appear in four tabs of big buttons at the bottom of the window. You have to go to the menu for less popular operations like rotating a photo.

Don't remove features or skimp on advanced features. Just don't shove them in the face of the end user. You'd be amazed how much friendlier a program can be if you simply consolidate and organize the features that exist into easy-to-use tabbed dialogs and wizards.

Developers are already hard at work making improvements like these, but many are not moving quickly enough. It's not fun to pump true usability into an application, but it is absolutely essential in order to take advantage of this window of opportunity.

5) Fill in the gaping holes in free software with more free software.

Free up some developer time to create robust equivalents for the remaining Windows applications for which there are no reasonable equivalents—iTunes and Quicken, to name two examples. Yes, there are personal finance applications, but developers need to catch up enough to compete with things like Quicken.

Java will be free, as in GPL, soon. All distros should start installing Java by default now, in anticipation of when most if not all of the JVM will be GPL—add to that a number of excellent Java applications now missing from distros.

These are but a few suggestions, but they are necessary and require urgent attention. Given enough time, people will migrate to Vista by default, simply because they upgrade their computers. If we can take advantage of the window of opportunity soon enough, we could see a Linux desktop revolution in the very near future. ■

Nicholas Petreley is Editor in Chief of *Linux Journal* and a former programmer, teacher, analyst and consultant who has been working with and writing about Linux for more than ten years.



Rackspace – Managed Hosting Backed by Fanatical Support™

Fast servers, secure data centers and maximum bandwidth are all well and good. In fact, we invest a lot of money in them every year. But we believe hosting enterprise class web sites and web applications takes more than technology. It takes Fanatical Support.

Fanatical Support isn't a clever slogan, but the day to day reality our customers experience working with us. It's how we have reimagined customer service to bring unprecedented responsiveness and value to everything we do for our customers. It starts the first time you talk with us. And it never ends.

Contact us to see how Fanatical Support works for you.

1.888.571.8976 or visit www.rackspace.com



Thanks for honoring us with the
2005 Linux Journal Readers' Choice Award for
"Favorite Web-Hosting Service"



The World's Fastest InfiniBand Switch

From a Company You've Trusted for 24 Years

Microway's **FasTree™** DDR InfiniBand switches run at 5GHz, twice as fast as the competition's SDR models. FasTree's non-blocking, flow-through architecture makes it possible to create 24 to 72 port modular fabrics which have lower latency than monolithic switches. They aggregate data modulo 24 instead of 12, improving nearest neighbor latency in fine grain problems and doubling the size of the largest three hop fat tree that can be built, from 288 to 576 ports. Larger fabrics can be created linking 576 port domains together.

Working with QLogic's InfiniPath InfiniBand Adapters, the number of hops required to move MPI messages between nodes is reduced, improving latency. The modular design makes them useful for SDR, DDR and future QDR InfiniBand fabrics, greatly extending their useful life. Please send email to fastree@microway.com to request our white paper entitled *Low Latency Modular Switches for InfiniBand*.



▲ 72 Port FasTree™ Configuration

Harness the power of 16 Opteron™ cores and 128 GB in 4U

Microway's **QuadPuter®** includes four or eight AMD dual core Opteron™ processors, 1350 Watt redundant power supply, and up to 8 redundant, hot swap hard drives—all in 4U. Dual core enables users to increase computing capacity without increasing power requirements, thereby providing the best performance per watt. Constructed with stainless steel, QuadPuter's RuggedRack™ architecture is designed to keep the processors and memory running cool and efficiently. Hard drives are cooled with external air and are front-mounted along with the power supply for easy access and removal. The RuggedRack™ with an 8-way motherboard, 8 drives, and up to 128 GB of memory is an excellent platform for power- and memory-hungry SMP applications.



Call us first at 508-746-7341 for quotes on clusters and storage solutions. Find testimonials and a list of satisfied customers at microway.com.

◀ QuadPuter® Navion™ with hot swap, redundant power and hard drives and four or eight dual core Opterons, offering the perfect balance between performance and density



508.746.7341 microway.com