

AUUGN

The Journal of AUUG Inc.
Volume 22 • Number 3
November 2001

Features:

Doing Damage with DocBook	14
AUUGN CD: Mandrake 8.1	16
Linux as a Replacement for Windows 2000	17
Interview with Ben Collins, DPL	27
Old Kernels Never Die	29
Rune for Linux	30
Testing a Firewall or Router	33
Caldera Open Linux Workstation 3.1	37
Wasabi Systems ships update to NetBSD	44
DVD Players for Linux	44
Using XML and XSLT	53
Logical Volume Managers	57
Psionic Portsentry, the defender of the ports	59
The State of Corporate IT: The case for Linux	62
Installing Debian on a 4MB RAM System	66
wxPython for Newbies	67
Debian Conference, Australia	72
AUUG2001: The Conference Photos	72
Shell Programming	78

News:

Public Notices	10
AUUG2002: Call for Papers	6
LUV Installfest 2001	12
AUUG Security Symposium 2001	86
AUUG: Chapter Meetings and Contact Details	90

Regulars:

President's Column	3
/var/spool/mail/auugn	4
My Home Network	10



AUUG Membership and General Correspondence

The AUUG Secretary
PO Box 366
Kensington NSW 2033
Telephone: 02 8824 9511
or 1800 625 655 (Toll-Free)
Facsimile: 02 8824 9522
Email: auug@auug.org.au

AUUG Management Committee

Email: auugexec@auug.org.au

President

David Purdue
iPlanet e-commerce solutions
The Tea House
Level 1
28 Clarendon Street
South Melbourne, Victoria, 3205
<David.Purdue@auug.org.au>

Vice-President

Michael Paddon
<Michael.Paddon@auug.org.au>

Secretary

Greg Lehey
IBM Australia
PO Box 460
Echunga, SA, 5153
<Greg.Lehey@auug.org.au>

Treasurer

Luigi Cantoni
Objective Management Pty Ltd
PO Box 51
North Perth WA 6906
<Luigi.Cantoni@auug.org.au>

Committee Members

Warren Toomey <warren.toomey@auug.org.au>
Sarah Bolderoff <sarah.bolderoff@auug.org.au>
Peter Gray <peter.gray@auug.org.au>
Conrad Parker <conrad.parker@auug.org.au>
Malcolm Caldwell <malcolm.caldwell@auug.org.au>

AUUG Business Manager

Elizabeth Carroll
PO Box 366
Kensington NSW 2033
<busmgr@auug.org.au>

Editorial

Con Zymaris
auugn@auug.org.au

I often begin issues of AUUGN with a snippet of recent news which relates to our segment of the IT industry. Now, what do I define as our segment of the IT industry? Well, to be blunt, and at the risk of inflaming opposing opinions, I define it as that area bounded by switched-on people doing interesting things with advanced computing platforms. Historically, these platforms have been Unix® and Unix-like, and TCP/IP-linked. Furthermore, these platforms, over the years, have been accruing (what I perceive to be) an increasing openness and vendor-independence. The culmination of this trend is the open source platforms and technologies running pretty much on everything from wristwatches through SPARC-stations and onto super-clusters.

It is interesting to view the link with the people who use these platforms; AUUG people and their ilk, ISOC-AU and SAGE-AU people have, over the past 20+ years, picked winners in Unix over proprietary OSes, TCP/IP over proprietary network protocols, the Internet over proprietary information infrastructure networks, and in more recent times, open source, multi-platform/vendor and POSIX-compliant environments and toolsets over proprietary, single-vendor ones. I think it's safe to say that the people who have chosen these paths have shown some foresight and have proven to be correct time-after-time.

Onto the recent news items. The US government has used the recent threats of terrorism as a welcome cache-prefetch with which to pipeline a number of IT-related policies which would have been politically impossible to even broach a year ago. Included in these are anti-privacy, anti-crypto and pro-content-control measures which are raising the ire of an increasing number of our colleagues overseas. While these activities are not happening here yet, this kind of governmental handiwork seems to spread.

Now, as a group of switched-on people within our industry, ones who have seen further and been proven correct on the merits and advantages of open and interoperable platforms and technologies, ones who firmly believe in the necessity of cryptographic technologies for ensuring individual rights to privacy, of systems security and open and publishable research therein, the members of AUUG have, I believe, an obligation to our industry specifically and to our community in general. This obligation is in advocacy and in the education of others and if possible, in the lobbying of the functionaries of government. If we, who often glimpse the future of technology, of its benefits and of the possibilities for its abuse, if don't do this, who will?

Cheers,

Con

Contribution Deadlines for AUUGN in 2002

Volume 23 • Number 1 – February 2002: **February 15th, 2002**

Volume 23 • Number 2 – May 2002: **May 15th, 2001**

AUUGN Editorial Committee

The AUUGN Editorial Committee can be reached by sending email to:
auugn@auug.org.au

Or to the following address:
AUUGN Editor
PO Box 366
Kensington NSW 2033

Editor:
Con Zymaris

Sub-Editors:
Mark Neely
Jerry Vochteloo

Public Relations and Marketing:
Elizabeth Carroll

AUUGN Submission Guidelines

Submission guidelines for AUUGN contributions can be obtained from the AUUG World Wide Web site at:

www.auug.org.au

Alternately, send email to the above correspondence address, requesting a copy.

AUUGN Back Issues

A variety of back issues of AUUGN are still available. For price and availability please contact the AUUG Secretariat, or write to:

AUUG Inc.
Back Issues Department
PO Box 366
Kensington NSW 2033

Conference Proceedings

A limited number of copies of the Conference Proceedings from previous AUUG Conferences are still available. Contact the AUUG Secretariat for details.

Mailing Lists

Enquiries regarding the purchase of the AUUGN mailing list should be directed to the AUUG Secretariat.

Disclaimer

Opinions expressed by the authors and reviewers are not necessarily those of AUUG Inc., its Journal, or its editorial committee.

Copyright Information

Copyright © 2001 AUUG Inc.

All rights reserved.

AUUGN is the journal of AUUG Inc., an organisation with the aim of promoting knowledge and understanding of Open Systems, including, but not restricted to, the UNIX® operating system, user interfaces, graphics, networking, programming and development environments and related standards.

Copyright without fee is permitted, provided that copies are made without modification, and are not made or distributed for commercial advantage.

President's Column

David Purdue
David.Purdue@auug.org.au

capacity *n.* (*pl. -ies*) power to contain, receive, experience or produce; maximum amount that can be contained etc.; mental power; position or function.
– *The Australian Little Oxford Dictionary*

At the time of writing two significant events for IT are about to occur: the Federal Election and the Melbourne Cup.

I am sure you will grant that these are both significant events – one of them the "race that stops the nation", the big chance that a bunch of nags have to show us what they are made of, and the other is a major horse race. But what is the connection to IT?

Well, this is the chance for the IT people at the Australian Electoral Commission and at Tabcorp to find out how good their capacity planning skills are.

The purpose of capacity planning is to ensure that your computing facility always has enough grunt to carry out the work required of it, and no more. The aim is to ensure that the old system runs out of grunt just at the time that the replacement system goes in to production.

Most of the science behind capacity planning was developed in the 1970's – the age of the mainframe. Good capacity planning was essential, because that next upgrade could cost in the order of \$1,000,000, and lead times for mainframe systems could be as long as 6 months. If you buy too soon, then you end up paying for capacity you don't need – a waste of a considerable amount of money. But order too late, and your data processing centre won't be able to get all the invoices out on time.

In the UNIX and especially the PC space of the 1990's and today, a lot of the science of capacity planning has been forgotten. At a time when the cost of computing power is continually dropping (CPUs get more powerful, RAM and disk capacities go up, but their cost declines), and delivery times are very short, the best strategy seems to be to wait until you need capacity and buy it then.

However not everything is so obvious. Recently, some computer scientists had a think about the implications of Moore's Law – which states that the speed of microprocessors doubles about every 18 months. They found that for problems that were going to take CPU time on the order of years, your best bet was to twiddle your thumbs for a while then buy your computer – you'd finish sooner.

There is a lot in the field of capacity planning that is applicable to today's computing environment. A lot of work has gone in to ensuring best use is made of the computer you have now, not just planning when the next one is needed.

In many ways, Tabcorp and the AEC have an easy task

of capacity planning – they both know when their peak load will be. In the case of the AEC, they have very good estimates of how many votes will be counted, for how many candidates, in how many polling booths. They know what information has to be transmitted to the national tally room, and where that information has to be sent afterwards. From this information they can predict, with some accuracy, their network and computing requirements on election night.

In any case, performance management is just one of the areas that will be covered at the next AUUG annual conference: "AUUG'2002 – Measure, Monitor, Control." I am pleased to announce that Adrian Close has accepted the job of Programme Committee Chair for the conference, and I thank him for volunteering.

You will find a call for papers in this issue of AUUGN, so now is the time to start thinking about what you can talk on, and get your abstract in.

/var/spool/mail/auugn

Editor: <auugn@auug.org.au>

What follows are some of the AUUG-related email exchanges which have crossed your editor's desk in recent times. If you want to contribute to the list, mail majordomo@tip.net.au with:

subscribe talk Your Name <your@email.com.au>

Date: Thu, 30 Aug 2001 09:57:22 +1000 (EST)

From: Damian Carvolth

<citecdjc@citecub.citec.qld.gov.au>

To: <talk@auug.org.au>

Subject: CD's in AUUG journal

Precedence: bulk

I figure since this list sees so much traffic...

I read with great interest the article in the current journal about the CD's.

I too had though they are a great idea, I've even used the Red Hat ones..

While I see people's interest in BSD I personally would prefer to get solaris 8 X86. The reason is Suns "media only" price is pretty exe. You can get BSD cheap as from numerous sources but since I mostly work on Solaris systems I'd like to+play with it at home...

This begs the question whether the solaris stuff can be included. Would sun allow it ?

And Staroffice sucks...sorry tried several versions. It just does...

Anyone care to perpetuate a religeous war ? :D

--

Damian Carvolth CITEC UNIX support.
damian.carvolth@citec.com.au
ph: 322 76334 f: 322 22684

Date: Thu, 30 Aug 2001 10:55:52 +1000

From: Greg Black <gjb@gbch.net>

To: damian.carvolth@citec.com.au

Cc: talk@auug.org.au

Subject: Re: CD's in AUUG journal

Damian Carvolth wrote:

| I read with great interest the article in the current
| journal about the CD's. | I too had though they are
| a great idea, I've even used the Red Hat ones..

I agree that they are a great idea. I haven't used any yet (since I already had the first sets), but I'm definitely going to give the OpenBSD set a try.

| While I see people's interest in BSD I personally

| would prefer to get solaris 8 X86. The reason is
| Suns "media only" price is pretty exe. You can get
| BSD cheap as from numerous sources but since I
| mostly work on Solaris systems I'd like to play with
| it at home...

I would certainly add any Solaris CDs to my coaster stack.

| And Staroffice sucks...sorry tried several versions. It
| just does...

Well, I was planning to push for StarOffice to be the next offering -- I'd like to give it a try.

Of course, if there's an alternative that does a better job, I'd be happy to try that instead (especially if it had the ability to correctly interpret those !@#%^& Word documents that people insist on sending to my wife).

Date: Thu, 30 Aug 2001 13:59:14 +1000 (EST)

From: David Purdue <davidp@canb.auug.org.au>

To: Greg Black <gjb@gbch.net>

cc: talk@auug.org.au

Subject: Re: CD's in AUUG journal

On Thu, 30 Aug 2001, Greg Black wrote:

>

> I agree that they are a great idea. I haven't used
> any yet (since I already had the first sets), but I'm
> definitely going to give the OpenBSD set a try.

Then please make sure you make a donation to the OpenBSD project - OpenBSD is funded primarily by CD sales and it was very generous of them to allow us to distribute the CD set.

DavidP

Date: Thu, 30 Aug 2001 17:48:46 +1000 (EST)

From: David J N Begley <d.begley@uws.edu.au>

To: <talk@auug.org.au>

Subject: Re: CD's in AUUG journal

Earlier today, Damian Carvolth wrote:

> And Staroffice sucks...sorry tried several versions. It
> just does...

Let's be honest - it's far from perfect, but the alternatives are pretty thin on the ground at present (not to say that there aren't alternatives, but those alternatives generally fit a niche market of users).

Sun StarOffice 5.2 has its problems (Sun Solaris SPARC version seems more prone to unexplained crashes than the Intel x86 Linux version), but at least it's fairly featureful and quite useful - I rely on the wordprocessor, spreadsheet and drawing components both for creating documents from scratch and also for reading MS Word and MS Excel documents sent to

me by colleagues.

No - converting documents to HTML, LaTeX or plain text is not sufficient in 99.9% of cases. Nor is generating those original documents in HTML or LaTeX (some documents are fine in HTML or LaTeX - others are not).

Yes, it can be slow/sluggish to start up - though once started it runs perfectly fine (he says, running StarOffice on a Pentium 133MHz with 96MB RAM); the startup time though is probably more a factor of C++ linking times, the same problem affecting KDE or other heavy C++ applications. The KDE developers believe they may have found a partial solution to this (certainly for Intel x86 users) and if it works, could easily be incorporated into the open source version of StarOffice - OpenOffice.

Maybe one day OpenOffice, KOffice or any other open source effort will be sufficiently "complete" to satisfy all our needs - but for now, many users can get a lot of value from StarOffice on UNIX, even if it falls short of meeting everyone's expectations.

Let's not write it off entirely - certainly not yet anyway. ;-)

Cheers..

dave

Date: Thu, 30 Aug 2001 21:53:01 +1000 (EST)
From: Glenn Satchell <Glenn.Satchell@uniq.com.au>
Subject: Re: CD's in AUUG journal
To: talk@auug.org.au

[... duplicated inline mail deleted...]

I like the CDs too. I actually installed all of them (one each quarter) on the spare disk I've got for the laptop, just to see what they were like.

If you download Solaris 8 it's 3 CDs, one "install" and two with all the packages, etc. there's an optional software CD with freeware ports.

When you download from www.sun.com you have to register as a user and nominate how many systems you will be installing. It's free to register and you can nominate as many systems in one go as you like (I think I said something like 5 new installs and 3 upgrades just to cover myself). Note that there are no license codes involved or any license enforcement during the install.

Or if you feel like paying for hard media, then you can get it for \$US45 plus shipping, which will probably be <\$150 Aus all up.

<http://www.sun.com/solaris/binaries/>

So, Sun _may_ allow inclusion of Solaris 8 CDs if we all promise to register on their web site if we're actually going to use them. I haven't checked this

with Sun and I don't know whether this would be acceptable or not. I suspect that their licensing scheme is to allow them to track approximately how many systems are installed.

As for Star Office. I use version 5.2 for Sparc quite a bit. It's heaps better than 5.1, and generally I have found it to be ok. If you use it's own staroffice format then it works very well, if a bit slow. Goes great on the 300MHz Ultra 10 w/ 256MB at the office. If you ask nicely Sun or your local reseller will give you a CD.

For converting Word to text I find strings to be your best friend, as long as you don't mind deleting all the font goo at the end.

regards,
-glenn

Glenn Satchell,
<mailto:glenn@uniq.com.au>
Senior Technical Consultant,
<http://www.ua.uniq.com.au/>
Uniq Advances Pty Ltd, phone: 0409-458-580
PO Box 70 Paddington NSW 2021

Date: Thu, 30 Aug 2001 11:22:25 +1000 (EST)
From: Damian Carvolth
<citecdjc@citecub.citec.qld.gov.au>
To: <talk@auug.org.au>
Subject: Word docs and staroffice

%s/^/> /g is a wonderful thing...

Anyway,

Regarding word docs at the end of this is an email from a wonderful lady friend of mine who cured the problem for me forever...


Staroffice does a good but not perfect job of converting. It also has a number of key problems. One i came across a while back was I built a database and reopening it it came up as a spreadsheet. I couldn't find any way to open it as+a database. This really annoyed me. Also it's taken it upon itself to rewrite html I opened in it without bothering to tell me. It's big, takes a long time to open and the interface has some problems and "lack of intuitiveness" IMO, the help also sucks.

The suggestions below may not suit you. You can try KDE2. I've got it on my linux box here at work and it seems reasonable.

There are many options as with all thing unix. The worst bit is finding time to find the one that suits you best...

Hope this helps...

http://www.onlamp.com/pub/a/bsd/2001/07/19/FreeBSD_Basics.html?page=2



Call for Papers: AUUG 2002 Theme: "Measure, Monitor, Control"

The AUUG Annual Conference will be held in Melbourne, Australia, on 4, 5 and 6 September 2002 (*subject to change*).

The Conference will be preceded by three days of tutorials, to be held on 1, 2 and 3 September 2002.

The Programme Committee invites proposals for papers and tutorials relating to:

- Cluster Computing
- Managing Distributed Networks
- Performance Management and Measurement
- Open Source Systems Administration Tools
- System and Application Monitoring
- Security in the Enterprise
- Technical aspects of Computing
- Networking in the Enterprise
- Business Experience and Case Studies
- Open Source projects
- Business cases for Open Source
- Technical aspects of Unix, Linux, and BSD variants
- Open Systems or other operating systems
- Computer Security
- Networking, Internet (including the World Wide Web)

Presentations may be given as tutorials, technical papers, or management studies. Technical papers are designed for those who need in-depth knowledge, whereas management studies present case studies of real-life experiences in the conference's fields of interest.

A written paper, for inclusion in the conference proceedings must accompany all presentations.

Speakers may select one of two presentation formats:

Technical presentation:

- A 30-minute talk, with 10 minutes for questions.

Management presentation:

- A 25-30 minute talk, with 10-15 minutes for questions (i.e. a total 40 minutes).

Panel sessions will also be timetabled in the conference and speakers should indicate their willingness to participate, and may like to suggest panel topics.

Tutorials, which may be of either a technical or management orientation, provide a more thorough presentation, of either a half-day or full-day duration.

Representing the largest Technical Computing event held in Australia, this conference offers an unparalleled opportunity to present your ideas and experiences to an audience with a major influence on the direction of Computing in Australia.



Call for Papers: AUUG 2002 Theme: "Measure, Monitor, Control"

Submission Guidelines:

Those proposing to submit papers should submit an extended abstract (1-3 pages) and a brief biography, and clearly indicate their preferred presentation format.

Those submitting tutorial proposals should submit an outline of the tutorial and a brief biography, and clearly indicate whether the tutorial is of half-day or full-day duration.

Speaker Incentives

Presenters of papers are afforded complimentary conference registration.

Tutorial presenters may select 25% of the profit of their session OR complimentary conference registration. Past experience suggests that a successful tutorial session of either duration can generate a reasonable return to the presenter.

Please note that with the GST changes to tax legislation we will be requiring the presentation of a tax invoice (which we will assist in producing) containing an ABN for your payment. If that is not provided then tax will have to be withheld from your payment.

Important Dates

Abstracts/Proposals Due	- 10 May 2002
Authors notified	- 7 June 2002
Final copy due	- 6 July 2002
Tutorials	- 1-3 September 2002
Conference	- 4-6 September 2002

Proposals should be sent to:

AUUG Inc.
PO Box 366
Kensington NSW 2033
AUSTRALIA

Email: auug2002prog@auug.org.au

Phone: 1800 625 655 or +61 2 8824 9511


Fax: +61 2 8824 9522

Please refer to the AUUG website for further information and up-to-date details:

<http://www.auug.org.au>

UNIX Instructors

Be the FIRST to learn new technologies



This is a rare opportunity to teach the next generation of UNIX gurus!! Become a key member of the premier global team of Solaris experts, and enhance your technical expertise by being among the first to learn the very latest UNIX technologies. Our client wants you to have as many career opportunities as possible – learning and training in other technologies, including JAVA and associated tools are options for you. Your enjoyment of sharing knowledge with experienced UNIX experts, combined with a strong systems administration or Shell programming background, will ensure your success in this challenging role. Training experience is not essential and you'll be instructing some 60 % of your time. Remuneration, benefits and career prospects are exceptional. Contact Andrew Parker for further information. (Ref. 0102-5)

PACT Global

Ground Floor, 41 McLaren Street, North Sydney NSW 2060 www.pactglobal.com.au
Tel: (02) 9954 4800 Fax: (02) 9954 4944 Email: mail@pactglobal.com.au

Public Notices

UPCOMING CONFERENCES

December 2-7

15th Systems Administration Conference (LISA 2001)

San Diego, CA

January 28-29, 2002

FAST - First Conference on File and Storage Technologies

Monterey, California

February 6 - 9, 2002

linux.conf.au

University of Queensland,
Brisbane, Australia

February 11-14, 2002

BSDCon 2002

Cathedral Hill Hotel
San Francisco, California

June 9-14, 2002

2002 USENIX Annual Technical Conference

FREENIX submissions deadline: November 12, 2001
General Session submissions deadline: November 19, 2001

Monterey Conference Center
Monterey, CA

USENIX IMPORTANT NOTICE ATA, PATRIOT, SSSCA

Your editor received the following important message from the USENIX Board of Directors. Please take the time to read it and visit the web-page listed. While this strictly involves legislation being discussed in the United States, this kind of legal baggage has the frequently annoying habit of wafting across the Pacific, and infecting Australia. If there is any way you can help, it's best to do so now, and help nip it in the bud.

The USENIX Board of Directors has decided to alert our membership that bills pending before the U.S. Congress or in committee appear to have a detrimental impact on computer professionals.

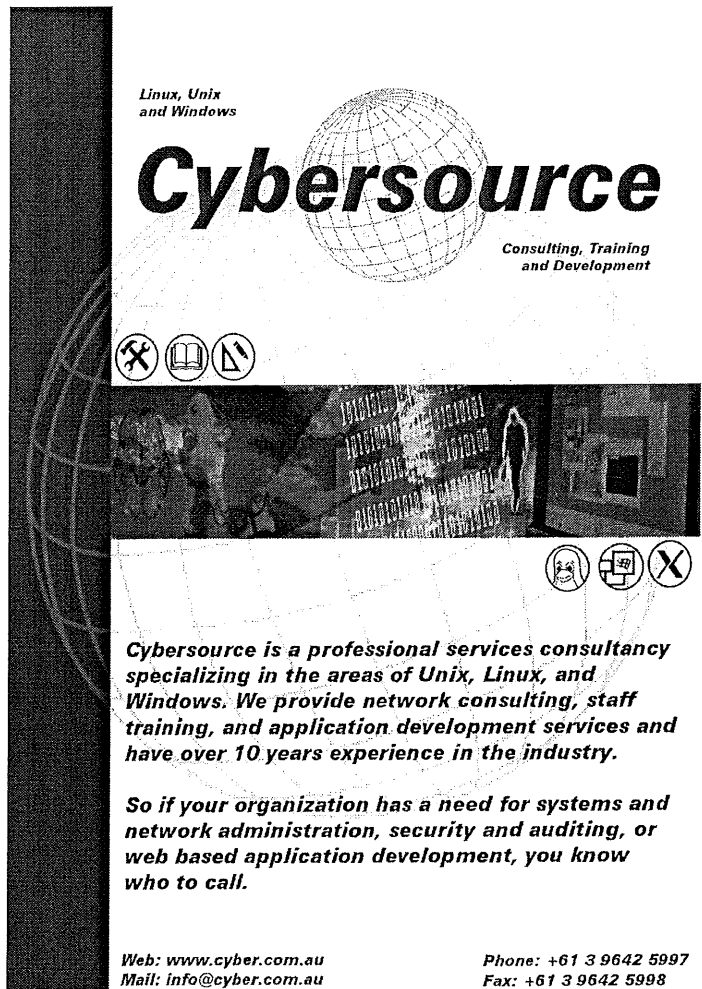
We are most concerned about aspects of two proposed bills, the Anti-Terrorism Act (ATA) and the Security Systems Standards and Certification Act (SSSCA), and how they interact with existing legislation such as the Computer Fraud and Abuse Act and the Digital Millennium Copyright Act. The ATA redefines virtually all computer crime as terrorism, enlarges the maximum penalty to life in prison without parole, allows broad pre-conviction seizures, and, furthermore, does this retroactively, removing the existing statute of limitations. The SSSCA essentially mandates copyright protection in all digital consumer devices and makes disabling or

avoidance of copyright mechanisms a felony offense.

More information about these and related issues can be found at:

<http://www.usenix.org/whatsnew/legislation.html>

Similar legislation is also being considered in other jurisdictions including Canada and some states. If you believe that provisions of these or similar acts are inappropriate we strongly encourage you to contact your elected representatives as soon as possible and register your opinions.



Linux, Unix
and Windows

Cybersource

Consulting, Training
and Development

⚙️ 📖 ⚠️

⚙️ 📖 ⚠️

Cybersource is a professional services consultancy specializing in the areas of Unix, Linux, and Windows. We provide network consulting, staff training, and application development services and have over 10 years experience in the industry.

So if your organization has a need for systems and network administration, security and auditing, or web based application development, you know who to call.

Web: www.cyber.com.au
Mail: info@cyber.com.au

Phone: +61 3 9642 5997
Fax: +61 3 9642 5998

My Home Network (November 2001)

By: Frank Crawford <frank@crawford.emu.id.au>

In the last issue I raved on about my new ADSL link and the fun and games with reconfiguring my whole external link to take advantage of it. However, I finished the article with one problem still outstanding: "How do external applications contact my server?". The problem is that my BigPond ADSL link supplies a dynamically allocated IP address, i.e. one that changes regularly, and in particular whenever the link is reset such as by a reboot.

This causes problems with such things as running my own mail server, web server, and ssh server, because each of these need to have a known IP address to connect, at least at the time of connection. The cost of a fixed IP address is considerably more, when it is available. Even more importantly, most home users don't require it, as they initiate connections, rather than act as the destination for a connection. (Note, don't consider this as a security feature as, while the addresses may not be published, they are available through such techniques as scanning network ranges, and other probes.)

There is one loophole in all this that allows connections to be established, and this relates to how sites are looked up. Rather than using an IP address, most sites are referred to by their name, which is then translated to an IP address by a DNS server. If this server is regularly updated with the current address, then lookups will return the correct address.

Of course, the problem with this is in the details. Items such as cached entries, timeouts and propagation delays all combine to make IP updating a tricky business. At the same time, dynamic updating of IP addresses (as this process is known) is beginning to be widely deployed. For example Microsoft Active Directory relies on it heavily, and it has been implemented in a number of DHCP servers. Unfortunately, all these examples are really applicable to an intranet, and do not easily scale to the Internet. In addition, such techniques as Dynamic Updates from a DHCP server are something that would need to be implemented by the ISP, not from the client end.

However, a number of mechanisms have been implemented to allow dynamic updates of IP addresses to work across the Internet. Invariably, these methods are specific to each service and tend to be web based. In their simplest form, each service has the facility to connect with a web browser and update the entry for the client host. What makes this possible is that the main purpose of updating the entry is to allow remote hosts to access the local client, it has nothing to do with the local client establishing a connection.

Because the local hosts IP address is allocated and

usable, it is simple for it to connect to a remote server (usually a web server) and update the entry recorded for their host with the current IP address. Once the address is updated in the DNS server, it is necessary for others on the Internet to use the latest address, rather than the old address. This is handled by ensuring that an address is not cached for any long period, with the cache timeout usually measured in seconds, not hours or days which you would set for a server with a fixed or static address.

This does have an impact on DNS performance, in that virtually every DNS lookup for your address must be referred back to a single server. If this was the case for a heavily access server, say Google's web server, then this would cause unacceptable delays and congestion, but for the sort of host that requires such a dynamic server, there are only a couple of lookups a day, and so it isn't a major effect.

So given that it was necessary for me to use such a service, my problem was to select a suitable service. There are a number, although I quickly limited my choices to two, ddns.nu, an Australian site, which made specific mention of supporting BigPond, and dyndns.org, a major international site. I should note, that there were a number of other suitable international sites, but dyndns.org offered all the tools and services I wanted.

The major differences between the two sites were:

- one was local and specifically had support for BigPond, while the other wasn't, and
- one was a major well organised site, with support for Linux and similar systems.

After some analysis I decided to go with dyndns.org as location wasn't as important as the support I could receive. So, to register I just connected to <http://www.dyndns.org> and added my host to one of their available domains (as 'crawford.dnsalias.org' in my case) via their web pages.

Obviously, manually updating the IP address every time it changes and doing it in a timely fashion is not practical, so it is necessary to automate the whole procedure. This was one of the reasons for selecting dyndns.org, there were a range of software to download. Of the ones available I chose 'ipcheck' (<http://ipcheck.sourceforge.net/>), a Python script that conforms to all the interface requirements of the dyndns.org site. One of these requirements is that you should not update the DNS site too frequently or unnecessarily (i.e., with the same address). Breaches of these requirements may cause your sites entry to be disabled until manually reenabled by one of the dyndns.org administrators.

Given that ipcheck conforms to all the requirements, it is recommended that it be run when the PPP link is first established, and also be run from cron occasionally, to catch any updates that it misses. In my case I found it easiest to put together a small

script that is called from either "/etc/ppp/ip-up.local" or from cron, every hour. The script looks like:

```
#!/bin/sh

IPCHECK=/usr/local/sbin/ipcheck
IPCHECK_DIR=/var/lib/ipcheck
IPCHECK_ACCT=/etc/ppp/ipcheck.acct

PPPOE_PID=/var/run/pppoe.conf-adsl.pid.pppd

if [ -z "$1" -a -n $PPPOE_PID ]
then
  IF_NUM=`cat $PPPOE_PID`
  PPIF=`grep "^$IF_NUM$" /var/run/ppp*.pid
/dev/null | sed
+'s:/var/run/(ppp.*)\.pid.*$:\1:'`
else
  PPIF=$1
fi

[ -f ${IPCHECK_DIR}/ipcheck.dat ] || \
  $IPCHECK -i ${PPIF:-ppp0} -d $IPCHECK_DIR
--makedat --acctfile
+$IPCHECK ACCT 2>&1 | \
  /usr/bin/logger -t ipcheck -s
$IPCHECK -i ${PPIF:-ppp0} -d $IPCHECK_DIR
--acctfile $IPCHECK ACCT 2>&1 | \
  /usr/bin/logger -t ipcheck -s
```

This script works out the appropriate interface, initialises the ipcheck database (if necessary) and then runs ipcheck to update my IP address with dyndns.org. If there is no change in the IP address, which there shouldn't be in most cases, then ipcheck doesn't forward the changes. It also logs all the output to syslog.

While this gets me a name that can be looked up via DNS, it is unfortunately, not the name I want to use. With everything on the Internet, there are a number of different possibilities, one of which was to find someone to host my domain in a fashion I could automatically update (something even dyndns.org do for a price) or have my domain entries point to the dyndns.org entry I had established. In this case I decided to take the second alternative.

As an MX record for my mail server was already being hosted on the AUUG DNS server, I asked if they would add a number of CNAMEs pointing to the host 'crawford.dnsalias.org'. Initially they raised a number of valid objections, which I've included here:

Stephen Rothwell wrote:

"I am not convinced that having mail delivered to an address that changes is sensible. Think about:

- 1) machine A looks up the MX record and gets your IP address
- 2) your link goes down
- 3) machine B's link comes up and gets your address
- 4) machine A deliver's some mail to your old address i.e. machine B now be one of the following could happen:
 - 1) machine B rejects your mail because it is not running an MTA - mail will be queued for later delivery
 - 2) machine B rejects your mail because it IS running an MTA and doesn't recognize the to

address - your mail will bounce.

- 3) you get luck and machine B is an open relay :-)
- 4) machine B accepts your mail (I am sure I could set up an MTA to do this) ..."

There was one additional problem I found, without having a second MX host that was available when my host was offline (for example, due to BigPond outages) then some mail would bounce. This is more because of poor MTAs, but it is still a fact of life.

After considering Stephen's comments, which were correct, I elected to still have DNS setup, generally in accordance with my original plan. In fact, there were a couple of technical issues that meant that there were some changes required, in particular, the MX for the domain name ("crawford.emu.id.au") points to my server, and a few CNAMEs were established for such things as "www.crawford.emu.id.au".

It is interesting to look back at Stephen's comments after a few months of running, to see how they stack up. In fact there have been a few times that when my dyndns.org entry was not updated, once due to an error in my update script (which meant I was online, but with a different address) and other times due my ADSL link being down.

The most common occurrence was mail was not accepted by who ever had my old IP address (i.e. case 1). This is not surprising, as the vast majority of home users will not have an MTA (or any other server) on their host, at least not intentionally. I have yet to hear of any mail sent to me being rejected due to connecting to another MTA (case 2), and haven't seen any mail coming from an open relay (case 3).

Finally, I guess that case 4 is a possibility, and for this reason, I have avoided receiving any mail that is confidential. If I was planning to, then I should be encrypting it anyway. :-)

So at this point I had a permanent connection (or at least as permanent as BigPond can make it), DNS pointing via a slightly round about method to my host and processes for keeping it up-to-date. I was ready to allow access from the outside world. From my last column, you should remember that I'd set up a number of special rules for the IPChains, and to these I added two additional ones, one for SMTP (i.e. to receive mail) and one for SSH (to connect remotely). As well, I'm considering opening up access to my Web server sometime, but not until it has some decent content!

There was one other port I found useful to open, AUTH, which is used by some MTAs as authentication mechanism. Access to this port is controlled by the 'identd' daemon, which can be configured to return various information.

So, now I have a permanent connection connected to my home, with access

from just about anywhere on the Internet. In

particular, I've found it very useful for work, where I can set up a number of processes to copy files back and forth, allowing me to work in either location. I used this very heavily to prepare the proceedings for AUUG2001. Digressing slightly, I heard an interesting presentation at the conference on the use of the Coda filesystem for keeping a laptop and a home system in sync. It would be interesting to see if I could use it here.

This permanent connection should make the world look rosy, right, ...wrong. Now I have to get very heavily into security issues, including things like virus protection. This is very much a work in progress issue, and I hope to document it in an upcoming column.

Now, on to other issues. I was thrilled to see Andrew McRae's home network in the last AUUGN. It got me thinking about how I plan to upgrade my network, and in general I agree with him that my next home network is likely to be a wireless network. However, I may wait a little while, to get a higher bandwidth link. Currently 802.11b has a maximum shared bandwidth of 11Mbit/s, but 802.11a has a bandwidth of 54Mbit/s, and, at least according to press releases from Intel, will cost less. These cards will become available next year and may be worth the wait.

Finally, I'm pleased to say that the AUUG Conference in Sydney went well. There were a number of interesting presentations, as well as a chance to get together with a wide range of experts. If you missed it, then pencil it into your diary for next year. I'd also recommend that you look into some of the other symposia that AUUG will be running around the country.

LUV Installfest: 2001

Linux Users of Victoria is proud to announce our participation in Installfest 2001, an Australia-wide Linux installfest effort. Our installfest will be held on the 24th of November.

The installfest is an easy way to get Linux installed on your computer. You can bring your PC to the installfest and:

- get help with installing Linux; or
- get help configuring your existing Linux system.

We will also be having ongoing demonstrations of Linux applications, and presentations throughout the day on what Linux is and what it can do for you. If you're curious about Linux and want to find out what it's all about, this is a perfect opportunity.

Details for the installfest are:

Date:

24th of November

Time:

10am - 5pm

Venue:

North Dining Hall
University of Melbourne
Melbourne Vic.

Cost:

FREE!, we hope

Other Installfest features: (Assuming we can get people to help arrange it)

free BBQ lunch and soft drink for all attendees between midday and 1pm.

free Caldera OpenLinux eDesktop 2.4 CD-ROM for those installing Caldera OpenLinux.

free Red Hat 6.2 CD-ROM for those installing Red Hat or upgrading an older Red Hat system.

hundreds of Linux door prizes and give-aways.

For those wishing to attend the installfest to have Linux installed or get help with their Linux installation, it is important that you register. Although we won't be turning people away on the day registered people will take preference.

We are also looking for LUV volunteers to help at the installfest. If you can help, fill out the volunteers registration.

Once you do this, please join the installfest mailing list by sending a mail to majordomo@luv.asn.au, with "subscribe luv-installfest (email address)" in the message body (subject does not matter).

MORE INFORMATION

<http://www.luv.asn.au/if/>

INFORMATION ON RUNNING AN INSTALLFEST

The 5 points on running an installfest can be found at www.linux.org.au/installfest/5points/. AUUGN would like to draw your attention to the

The 5 points on running an installfest can be found at www.linux.org.au/installfest/5points/

AUUG Corporate Members

as at 1 July 2001

- Andersen Consulting
- ANSTO
- Aurema Pty Ltd
- Australian Bureau of Statistics
- Australian Industry Group
- Australian Taxation Office
- Australian Water Technologies P/L
- BHP Information Technology
- British Aerospace Australia
- Bureau of Meteorology
- C.I.S.R.A.
- Cape Grim B.A.P.S
- Central Queensland University
- Central Sydney Area Health Service
- Centrelink
- CITEC
- Commercial Dynamics
- Commonwealth Steel Company
- Computer Science, Australian Defence Force Academy
- Computing Services, Dept Premier & Cabinet
- Corinthian Industries (Holdings) Pty Ltd
- Corporate Express Australia Limited
- Crane Distribution Limited
- CSC Australia Pty. Ltd.
- CSIRO Manufacturing Science and Technology
- Curtin University of Technology
- Cyberscience Corporation Pty. Ltd.
- Cybersource Pty. Ltd.
- Daimler Chrysler Australia - Pacific
- Dawn Technologies
- Deakin University
- Department of Defence
- Department of Land & Water Conservation
- Energex
- eSec Limited
- Everything Linux & Linux Help
- Fulcrum Consulting Group
- Fulcrum Consulting Group
- G.James Australia Pty. Ltd.
- IP Australia
- IT Services Centre, ADFA
- Land and Property Information, NSW
- LPINSW
- Macquarie University
- Mercantile Mutual Holdings
- Motorola Australia Software Centre
- Multibase WebAustralis Pty Limited
- Museum Victoria
- Namadgi Systems Pty Ltd
- Nokia Australia
- NSW National Parks & Wildlife Service
- NSW Public Works & Services, Information Services
- Peter Harding & Associates Pty. Ltd.
- Qantas Information Technology
- Rinbina Pty. Ltd.
- Security Mailing Services Pty Ltd
- Snowy Mountains Authority
- St. John of God Health Care Inc.
- St. Vincent's Private Hospital
- Stallion Technologies Pty. Ltd.
- Standards Australia
- State Library of Victoria
- TAB Queensland Limited
- The University of Western Australia
- Thiess Contractors Pty Ltd
- Tower Technology Pty. Ltd.
- University of Melbourne
- University of New South Wales
- University of Sydney
- University of Technology, Sydney
- Victoria University of Technology
- Westrail
- Workcover Queensland

Doing Damage with DocBook

Author: Paul Tatham <paul@sympatico.ca>

INTRODUCTION

This is not a tutorial on how to use DocBook; there are plenty of excellent ones available such as this one by David Ruggie, Mark Galassi and Eric Bischoff. Rather, it is an outline of things not to do from the "don't make the mistakes I made department"; a kind of "do as I say, not as I did" piece. Those who are thinking about writing documentation, especially for open source projects, may find it useful to learn the many ways to screw things up, just as I did. That way they will know what pitfalls to avoid. But first, a quick description of DocBook.

WHAT'S DOCBOOK?

DocBook, currently maintained by OASIS, is a markup language described by a document type definition for SGML or XML. A couple of years ago, if someone had told me that my eyes would have glazed over and I would have asked for a glass of water. Let's see if we can make it easier. DocBook uses a set of tags to describe the structure and content of a document, somewhat reminiscent of HTML. For example, a paragraph in DocBook might look like this:

```
<para>It will take 5 hours to get through Spike Milligan's Puckoon, 2 to read and 3 to laugh.</para>
```

In fact, HTML is really another document type definition. While you may be concerned with format, or how your web page will look when displayed in a browser with HTML, with DocBook you are concerned with describing your document semantically. Format is only applied when the document is processed and converted into a form suitable for publication. DocBook also has a great many more tags (the proper term is named elements) available than HTML. Its main purpose is for producing books and articles on computer hardware and software. But unlike HTML, you can maintain one document for output to several formats including HTML, Postscript, pdf, text, rich text format, and so on. This means you can publish in almost any medium from one source - it's portable. What about SGML and XML? These are really meta-languages that describe how markup languages like DocBook or HTML are defined. As far as Docbook is concerned, there are few differences between SGML and XML. Examples: with SGML you can use tag minimization like `<para>contents</>` and you can ignore case - `<para>` is the same as `<PARA>`. But with XML all tags must be lower case and tag minimization is not permitted. Your best bet is to follow the XML convention, even if you are using SGML. That way you can change your document from SGML to XML by changing one or two lines at the

beginning. Beyond that, if documentation creation is what you want to do, follow the advice of Michael Smith when he says "don't learn XML", and make sure not to learn SGML as well. DocBook is all you really want to know.

What's in a DocBook Document?

The Document Type Declaration All SGML DocBook documents start with something like this:

```
<!DOCTYPE article PUBLIC "-//OASIS//DTD DocBook V4.1//EN">
```

This tells whatever software is used to process the document that the Document Type Definition being used is DocBook, and that the DTD can be found with the public identifier, namely OASIS version 4.1. It also tells us that the document is an article; the lingua franca of DocBook says that article is the root element, kind of like `<body>` in HTML. The root element can also be book. If your document is XML it will need an XML Declaration like this:

```
<?xml version="1.0" ?>
```

```
<!DOCTYPE book PUBLIC "-//Norman Walsh//DTD DocBk XML V3.1.4//EN"
```

```
"file:///usr/local/docbook/xml/3.1.4/db3xml.dtd">
```

The Document Type Declaration also needs to show where the DTD can be found on the local system, rather than just giving the public identifier, which in this case is the Norman Walsh DTD. But apart from that, the rest of the document can be the same for both XML and SGML, as long as you follow the more restrictive method of markup (no tag minimization, only lower case, and a few minor things you can ignore for now). Which should you use, SGML or XML? If documentation is the object, just use SGML. XML is really just a new version of SGML with stricter rules to make it easier to create software that is XML aware, and this mainly in the area of structured information that is to be used over a network. But if you are writing books or articles, it really makes little difference.

Elements Elements are just the tags along with whatever content is included between them. There are lots of them though; try to learn the important ones and look up others as you need them. Nobody knows all of them so don't bother trying. DocBook best practices says you should use as much markup as possible; it might not matter if a given tag is present for one output format but it might for another. For example:

```
<para>Mandrake Linux is easy to install.</para>
```

might be better if it was written like this:

```
<para><productname>Mandrake Linux</productname> is easy to install.</para>
```

It won't make any difference when converted to HTML but it might when processed into Postscript for publication, depending on the stylesheet that's used. Elements can also have attributes like this:

```
<chapter label="1" id="Intro" xreflabel="Introduction">
```

Attributes are bits of other information that DocBook allows; amongst other things, they let you set up

hyperlinks in your documents.

ENTITIES

Entities are really just names for chunks of information. You can give a name to a chunk and then just use the name to refer to the whole thing. Actually there are two kinds of entities. General entities can be internal to DocBook and you don't have to declare them. These are things like the "&", ">" and "<" symbols which look like "&", ">" and "<"; in a document. In effect, these are special characters in DocBook that are escaped by making them entities. You can also have external entities that are declared in the prologue like this:

```
<!DOCTYPE article PUBLIC "-//OASIS//DTD DocBook
V4.1//EN" [
  <!ENTITY main SYSTEM main.jpg NDATA JPG>]>
```

This entity might be an image of a screenshot that you want to put in your document. Now you can just write `<graphic entityref="main"></graphic>` and DocBook will know that the `main.jpg` belongs there and it will be inserted when it is processed. It's useful when instances of things that appear in your document need to be changed, for example, things that depend on version numbers. DocBook also uses entities for marking sections of a document. Marked sections let you have different versions of the same document using "INCLUDE" and "IGNORE" parameters. Clever and useful.

INSTALLATION

Okay, you want to get started, but first you need to install DocBook. This is really easy; all you have to do is use a Linux distribution that installs everything you need by default. Many already do. But you want to install it yourself - that's easy too - don't; unless you really have nothing better to do. Because unless you want to mess around with OpenJade, the DocBook DTD, the DocBook DSSSL, tools like SgmlTools-Lite or Cygnus DocBook tools, ISO8879 entities, setting a bunch of environment variables, and otherwise fiddling with stuff; installing everything you need for DocBook is a waste of time. If I haven't persuaded you then try the excellent DocBook Install mini-HOWTO by Robert Easter. It's only 28 fun-filled, easy-to-read pages, and it's great in moments of desperation.

CREATING DOCUMENTS

So your machine is now DocBook ready and so are you. Now what? Well frankly you don't have many good choices for creating your documents. There are a few commercial products like epcEdit that, while quite good, can come with a hefty price tag. I tried a time-limited copy of epcEdit and found it was a real treat for new documents; it inserts tags for you and only lets you choose ones that are valid. But opening docs created elsewhere puts you in text mode without even tag highlighting. If there was a way around this I didn't find it. If you want to stick to open source tools

you'll have to use your favourite editor. VIM provides tag highlighting, but Nedit is also good. I'm told that the PSGML major mode for Emacs is the best of the bunch but for me, there aren't enough "fs" in "difficult" to describe Emacs. Beyond playing dunnet, forget it. Using an editor means you will have to type all the tags out yourself; you'll have to learn some of the rules of structure and hierarchy for DocBook and you'll have to validate your documents. Do this latter often; don't try to complete a large document and wait until the end to look for tag errors. The simplest thing is simply to process your `document.sgml` with either DocBook Tools or SgmlTools-lite. Let's say you want to produce HTML; here's what you type at the command prompt:

```
Bash-2.04$ db2html document.sgml
```

or

```
Bash-2.04$ sgmltools -b html document.sgml
```

The processor will try to make an HTML document from your input, but it will not succeed if there too many markup errors. You will need to use the error reporting displayed in the console to make the corrections to the line numbers where the errors are reported to have occurred. Of course it's easier if there are fewer errors. One of the reasons to use DocBook is to spend time thinking about structure and content rather than format. It's a good idea to plan your document before you write. Decide how it will be structured and start adding markup for the structure. What you are doing is creating a template that you will fill in afterwards with content. I can tell you that starting with content and adding structure later doesn't work. If you are planning to write for the Linux Documentation Project you will want to read the Author's Guide before you begin. You should also check the LDP resources page where you will find stylesheets and templates for your markup. Finally, you will find a wealth of information on Robin Cover's XML Cover Pages. It says XML but DocBook and SGML figure prominently. And after reading my complaints you will find it a welcome change. Have fun with DocBook!

BROWSE THE ARCHIVES OF THE VOICE OF THE BEHOLDIN'

If you have any concerns, questions or comments about this column please contact me at paul@mug.ca. I do not promise a reply but I will read any mail I receive.

ABOUT THE AUTHOR

Paul Tatham is a user and advocate of Free Software in general and GNU/Linux in particular. He discovered Linux in early 1999 after finding a need for an alternative, and has been using it ever since. He also began observing the growth and changes to the platform and decided to join the community to

contribute where he could. Although he has at times tried his hand at writing code he is now no longer permitted to do so. Paul lives with his wife and family in Rosemere, Quebec, the most beautiful town in Canada. He is also a member of MLUG, the Montreal Linux User Group.

This article is re-printed with permission. The originals can be found at:

<http://www.mlug.ca/paul/stories/column14.html>

AUUGN CD: Mandrake

8.1

This issues CD-set is the latest from Mandrake.

The following is some information from Mandrake's website as to what is included with this version

- **MandrakeFirstTime** allows new users to configure a window manager and email client (at present KMail and Netscape.)
- **DiskDrake** (disk partitioner) now supports network filesystems (Samba and NFS), as well as Windows network browsing.
- **Draknet** (network configuration tool): Enjoy a thoroughly reworked version.
- Improved Mime Type management
- Updated **URPMI** (package installer) and Software Manager
- The reworked Mandrake Control Center provides a convenient embedded root console as well as many new utilities such as Logdrake, a graphical frontend to the system's log files, and Drakinst, which allows easy creation of an auto-install floppy.
- Improved **HardDrake** (Hardware manager)
- Fresh version of **PrinterDrake** to get the most out of your printer.
- New Menus (usable via Menudrake) categorize applications by task, thereby making it easier for new Linux users to accomplish routine day-to-day tasks.
- Support for the Euro in few applications
- Enterprise Kernel 2.4.8: This special edition of the kernel provides native support for Hi Memory (>1024 MB) and SMP (Symmetric Multi Processing)
- Support for Journalized FileSystems EXT3, ReiserFS, XFS and JFS ensure filesystem integrity
- Special version of SAMBA: combined with XFS and our latest kernel, Samba 2.2.1.a allows Windows file sharing with NT-like ACLs (Access control lists).
- PHP Groupware: A full-featured groupware suite accessible from a secure web browser.

INSTALLATION AND CONFIGURATION

- Server configuration wizards: Many configuration wizards facilitate the configuring of main server functions (Web Mail, Firewall, File & Printer sharing,...)
- Mandrake Server CD: A complete server-oriented distribution featuring the best Open Source server software combined with Mandrake's easy to use installer.
- Web-based administration with Webmin: Configure DNS, Samba, NFS, local/remote file systems and more from any computer with a secure browser connection.
- URPMI: Automated software management systems URPMI (and APT) provide a hassle-free way to install and upgrade applications.
- Auto-Install tool: Duplicate servers and workstations effortlessly with the new "auto install" feature. User interfaces
- KDE 2.2.1 includes the dramatically improved KOffice 1.1 office suite
- GNOME 1.4.1, with
- Evolution 1.0 Beta 3: enjoy this newest version of the ultimate tool for personal and workgroup information management.
- Nautilus 1.04: the file manager for the GNOME desktop.

DESKTOP AND MULTIMEDIA SOFTWARE

- Grio500: synchronize your desktop with the Rio 500 MP3 player
- Mozilla 0.9.4: browse the web and try the new communication module
- XMMS 1.2.5: edit and manage MP3 files
- GIMP 1.2.2: create and manipulate photos with this powerful graphics software
- Gphoto2: manage all your digital photos
- Grip 2.96: burn you favorite CDs
- KOffice 1.1: perform all your office tasks
- Galeon 0.12.1: try this unusual browser for a new kind of browsing experience
- Gnomemeeting 0.11: Share good times with family & friends with this full-featured video conferencing software

Linux as a Replacement for Windows 2000

Author: Rob Valliere <rob@debianplanet.org>

THE BOTTOM LINE

This review focused on Red Hat Linux 7.1 from a business user's view and attempted to answer my client's question "Can Linux be used as a replacement for Windows 2000". After an intensive hands-on Linux project lasting several months, I was able to provide my client with a pertinent answer to this question.

- Based on a solid Linux business plan, my client decided that Linux was a good investment for most of the company's employees, but will retain a few key Windows computers in the office for specific applications. Connecting Linux and Windows computers on the same network, and sharing Office 97/2000 files between the different OSs, is easy and works well, thanks to Red Hat Linux and Star Office. The primary reason for this decision was a \$10,000 saving on his IT budget! For support reasons, he decided to use a local Linux supplier for installation and configuration work.

For other business users, here is my advice:

- For whatever reason, if you want to get started with Linux, or implement an office network of Linux and Windows clients, Red Hat Linux 7.1 is very much the way to go. Red Hat Linux 7.1 can be used as an alternative to Windows 2000! You will be stunned by the bang for the buck that Linux bundled free "open source" software offers. Red Hat Linux is a complete server (LAN, Web, SQL) and offers excellent desktop applications for corporate users. Each business user will need to look at the benefits of using Linux (no software licensing fees, fewer hardware upgrades, many good applications) and the costs (installation, configuration, upgrades, training, support). If the benefits outweigh the costs, then you have an affordable viable desktop alternative to Windows.

Red Hat Linux 7.1 is a great package: it provides a smooth installation, has many good applications and is an especially attractive option for small to medium-size offices. However, many business users will want to purchase a computer with Linux pre-installed, or at least use a Linux support professional for post-installation configurations, maintenance and upgrades, just as you currently do for Windows.

THE BACKGROUND

A client asked me whether Linux could be installed on

existing office hardware, and whether it was an alternative to Windows 2000. Considering the absence of software licensing fees with Linux and the ability to run Linux on limited hardware resources, the answer to this question may be of interest to other Windows 2000 business users. I was certainly interested.

He wanted to reduce his IT investments, from the expensive and 'never-ending' software and hardware upgrade cycles with Microsoft products, to a more practical and cost-effective strategy. He stated that reducing his IT budget, without compromising the applications, would be the deciding issue. My client stressed that if Linux was selected as the alternative OS, he "was not prepared to waste his time fighting IT fires; he wanted applications that worked with minimal fuss".

There were several factors relevant to this question, such as what applications were used in Windows, what older hardware was available and what kind of installation would be required (automatic or some what manual). My client provided the following list of requirements:

- Install on existing office hardware that is too slow for Windows 2000
- Offer an easy installation, similar to Windows 2000
- Provide office file and printer sharing between Linux and Windows computers
- Obtain technical support for Linux if and when needed
- Availability of good office applications, similar to the following Windows programs:

Application Type	Windows 2000 Application
Office 97	Word, Excel, PowerPoint, Access, Outlook
Internet	Internet Explorer 5, WS_FTP Pro 6
Security	ZoneAlarm Pro, Norton Anti-virus 2001
Publishing	Acrobat Writer 4, Adaptec Easy CD Creator 4
Web Graphics	ACDSee 3, Fireworks 4
Utilities	Acrobat Reader 5, WinZip 8, WinAmp 2.7

While the above software is "off-the-shelf", the office staff was dependant on a custom MS Access database application. We needed to find an alternative that provided the database backend and administrative tools required to maintain the application.

Other business users may have different requirements, for example, some would be willing to purchase new computers with Linux pre-installed. This would ensure there were no installation and configuration hassles. And many users may want different applications suited to their business needs.

MY BACKGROUND.

It is important that readers of this review understand my biases, which are in the MS DOS/Windows world. Further, my attitude on using computer technology comes very much from a practical, business view. I have read many times that Linux advocate's only use the Gnome GUI (KDE is never considered as an option), and believe that "Linux is not for everyone, blah, blah, blah". My attitude is that if Linux is user-friendly and works, then savvy business users will adopt it. And as most businesses use Windows, users will want a familiar desktop and will find KDE as good as, or even better than, Windows.

LINUX ON THE DESKTOP – A SUMMARY

The following is a summary of my review of Red Hat Linux 7.1, based on my client's requirements, and on my own experiences after using Linux daily for over 3 months. The Linux project included 6 installations, hardware, software and network configurations, testing, upgrading, problem solving and usage.

Install on existing office hardware

Yes. The Linux GUI environment can run on aging Pentium computers with 64 MB or more of RAM, but do your homework first to reduce some potential hardware configuration challenges.

Offer an easy installation

Yes. Excellent installation wizard facility, including on-screen help, 'Plug-n-Play' and Package Group selections. But all is not perfect. The biggest hurdle may be the configuration of your graphic hardware. Purchasing new computers with Linux pre-installed solves this headache.

Provide office file and printer sharing

Yes. Very good network file and printer sharing using Samba. Its best to get a professional to implement so server set-up, security, disaster prevention, and disaster recovery are guaranteed.

Obtain technical support

Yes. There is very good documentation and Web support forums available, but all this takes your time and effort. If you do not have in-house resources, you will need to find a local Linux supplier. If you want support from Red Hat directly, you will have to pay.

Availability of good office applications

Yes, with a few compromises. The KDE graphical desktop in Linux is excellent and functions like Windows 2000. Further, the vast majority of Linux applications included with Red Hat, for example Star Office, are available free of charge to end users. Here is a summary of the applications with my rating on how they compare:

UPDATED APPLICATION SUMMARY (

An updated summary of the applications with my rating on how they compare. Updated Linux applications are in **bold**. Note that an Outlook Express alternative was not included in the original review.

Alternative Rating:	Excellent	Capable	Limited
Application	Windows 2000		Free Red Hat Linux 7.1 Alternative
Word Processor	Word 2000		StarWriter: excellent Word Processor.
Spreadsheet	Excel 2000		StarCalc: very capable Spreadsheet.
Presentation	PowerPoint 2000		StarImpress: impressive presentation program.
Database	Access 2000		LAMP intranet using a PHP/MySQL application.
PIM with E-mail	Outlook 2000		StarSchedule: reasonable PIM. StarMail: separate E-mail/Address book. Not recommended
POP3 E-mail	Outlook Express		Sylpheed 0.6.2-1: Excellent Email client.
Browser	Internet Explorer 5		Mozilla 0.9.4 (an updated 8 MB download)
FTP	WS_FTP 6		gFTP 2.0.7: good ftp program with similar interface.
Firewall	ZoneAlarm Pro		Firestarter 0.7.1: Great GUI front-end for Linux Firewalls.
Anti-Virus	Norton Anti-virus		Not required with Linux, according to the experts.
PDF Writer	Acrobat Writer		StarWriter configured to print PDF directly.
CD-R Recorder	Easy CD Creator		Gtoaster 1.0Beta2: Excellent GUI for burning CDs.
Image Viewer	ACDSee 3		GQview 0.8.1: Very good image viewer. Or Konqueror.
Image Editor	Fireworks 4		Gimp 1.2: like PhotoShop! Great for Web graphics.
PDF Viewer	Acrobat Reader 4		Acrobat Reader 4 or xpdf 0.92: take your pick, both work.
Archive Utility	WinZip 8		Archiver (ark 1.9): excellent archiving program.
MP3 player	WinAmp 2.7		XMMS 1.2.4: like WinAmp, or KDE MultiMedia Player.

To find an alternative to the MS Access database

application, we decided to develop a prototype of an intranet application using Apache Web server, MySQL database and the PHP scripting facility. This Linux solution was not only faster and more powerful than the Access application, the software was available at no charge.

LINUX ON THE DESKTOP – THE DETAILS

While the summary was primarily aimed at providing my client with a concise answer to the question "Can Linux be used as a replacement for Windows 2000", the details present some fine points of the business plan and on how my client made his decision to use Linux and StarOffice as alternatives to Microsoft products.

Further, my recent Red Hat Linux experiences are outlined and where appropriate, my opinions provided. With 20 years in the MS DOS/Windows world but new to Linux, I was very keen to evaluate the LAMP Web development facility and if feasible, 'resurrect' my favorite 3 year old notebook.

TRIMMING \$10,000 OFF THE IT BUDGET

My client's LAN consisted of 25 computers: a Windows NT server and 24 Windows NT workstations using Office 97. The workstations had 128MB of RAM, okay for NT but not recommended for Windows 2000. The costs to upgrade this office to use Windows 2000 and Office 2000 were significant.

When it became clear that Linux could provide most of my clients requirements, especially the intranet database application and almost US\$10,000 could be trimmed off the IT budget, it was an easy business decision to implement Red Hat Linux as a replacement for Windows 2000. Linux would be deployed on the server for file, print, database and Web services, and on 20 workstations, it would replace Windows NT as the desktop OS. Here are the Windows 2000 Upgrade and Red Hat Linux Alternative IT budgets comparisons:

Windows 2000 Upgrade:

Workstation:	Qty	US\$	Total US\$	\$15,060
RAM upgrade	24	\$40	\$960	
Win2K upgrade	24	\$220	\$5,280	
Office2K Pro upgrade	24	\$330	\$7,920	
Server:				
Win2K upgrade w/25 CALs	1	\$900	\$900	

Linux Red Hat Alternative:

Workstation:	Qty	US\$	Total US\$	\$5,160
RAM upgrade	4	\$40	\$160	
Win2K upgrade	4	\$220	\$880	
Office2K Pro upgrade	4	\$330	\$1,320	
StarOffice	20	\$0	\$0	
Server:				
RedHat 7.1 Prof Server + shipping	1	\$300	\$300	
LAN installation, setup, training, SQL conversions, documentation		\$2,500	\$2,500	

Note: this budget comparison only included Microsoft upgrade fees, and not the software licensing fees for other programs used on Windows: for example, Acrobat Writer and Fireworks.

SELECTING A DISTRIBUTION

Not being an experienced Linux user, my first task was to select a Linux distribution. The top providers are: Mandrake, Red Hat, TurboLinux, VA Linux, SuSE, Debian and Caldera. [Thanks to several readers, Mandrake is now included.] I have not included Corel Linux, as its future will likely be influenced by one of its shareholders –Microsoft, who regards "open source" software as "a cancer".

Each distribution has its own setup program and range of applications. I selected Red Hat Linux as its distribution is aimed at business users and includes an excellent installation program and a full array of useful Linux tools, including both workstation and server applications. The latest release is 7.1 and their top of the line product includes 9 CDs full of software and documentation.

SELECTING THE HARDWARE

I selected 2 older systems from my office for testing. The Red Hat Linux 7.1 Professional Server package was installed, configured and tested on the following hardware:

- 2 year-old **Asus P2B-F desktop** (Pentium III 500MHz with 256 MB). Windows 2000 operates well on this hardware. There were 4 NTFS partitions: I left Windows on the NTFS drive C:, but converted the other 3 data drives to FAT32 using PartitionMagic 6.0. Linux was installed on the remaining free 6 GB. I used Linux's Lilo boot software to handle the multi-booting. Linux could access the 3 FAT32 data partitions, but not the NTFS partition.
- 3 year-old **Gateway Solo 9100 notebook** (Pentium II 300MHz with 96 MB). The RAM on this notebook was insufficient to run Windows 2000. Linux was installed on this notebook as the only OS.

All computers were connected to any existing Windows 2000 SP1 TCP/IP network.

Notes on Windows 2000.

Why compare Linux with Windows 2000 and not Windows 98/Me? Linux and Windows 2000 can be compared as both offer business users the security and reliability needed with office LANs. Both can be used as workstations or servers. Windows 98/Me, on the other hand, is much less secure, much less reliable and was never designed as a server.

What are the Windows 2000 Hardware

Requirements? Window 2000 demands a fast CPU (ideally starting with Pentium III) and a minimum of 256 MB of RAM. While Windows 2000 can run in 128 MB of RAM, when Office 2000 applications are loaded, it slows considerably.

LINUX HARDWARE REQUIREMENTS

These observations were based on a Pentium III 500 MHz desktop computer with dual OSs installed. Both Windows 2000 and Red Hat Linux 7.1 ran multiple services: web server, SQL database server, firewall, etc. The table clearly shows the significant differences in how memory is used in Windows and in Linux.

	RAM used after boot-up		RAM used after Star Office 5.2 or Office 2000 (Word, Excel, Outlook) loaded	
	Initial RAM Used	Swap file Used	RAM Used	Swap file Used
Red Hat Linux 7.1	45 MB	0 MB	70 MB	0 MB
Windows 2000 SP1	140 MB	35 MB	170 MB	35 MB

Linux can use existing office hardware as it requires considerably less hardware resources than Windows 2000! Linux is very efficient at memory usage and can run many programs, including Star Office, in 128MB without using a swap file. Without Star Office, runs well using a Pentium II with 96 MB of RAM. The conclusion: Linux runs fine in less than half the RAM that Windows 2000 requires, and does not require the fastest CPU on the market.

RED HAT HARDWARE COMPATIBILITY

This applies specifically to Red Hat 7.1's ability to automatically detect hardware. Hardware configuration was another matter, especially with the video cards. All of the following hardware was correctly detected by Red Hat 7.1, except for the 3 year old notebook's sound card and CompactFlash adapter:

Asus P2B-F Desktop

- Pentium III 500 MHz with 256 MB RAM
- Dual 13 GB Hard drives
- Drives: LS-120 and HP CDWriter+ 8100
- 15 " ADi ProVista monitor (1024x768 at 16-bit color)
- Asus AGP-V3100 graphics card (S3 Savage3D with 8MB RAM)
- Creative Sound Blaster Live! sound card
- D-Link DFE-530TX PCI network card
- Standard built-in facilities: Multi-I/O ports including USB
- Additional hardware
- Diamond SupraExpress 56e-PC V.90 PnP modem
- Epson USB Stylus Color 740 printer
- MS USB IntelliPoint Optical Wheel Mouse

Gateway Solo 9100 Notebook

- Pentium II 300 MHz with 96 MB RAM
- 6.4 GB Hard drive
- Combo Drive: Matsushita SR-8171 CD/DVD and LS-120
- 14 " LCD display (1024x768 at 16-bit color)
- Trident Cyber 9397 video chipset (4MB RAM)
- Yamaha OPL3sa sound chipset
- 2 Type II PCMCIA (Cirrus Logic PD6832 CardBus chip set)
- Standard built-in facilities: Multi-I/O ports including USB and
- Touchpad
- Additional hardware
- Xircom Credit Card Ethernet 10/100 + Modem 56 (CEM56-100)
- Adapter Card
- Kingston CompactFlash Adapter Card with CASIO 8 MB CompactFlash
- Card
- MS PS/2 IntelliMouse

Red Hat's support for hardware is very good, but not as good as Windows 2000. Hardware support is not an issue if you purchase a computer with pre-installed Linux.

RED HAT INSTALLATION

The Red Hat 7.1 installation was easy and almost as automated as Windows 2000. You simply boot the computer using the first Red Hat CD-ROM disk. There are good help screens, hardware is automatically detected, automatic partitioning is available, and you can choose from several Package Groups. However, and based on my experiences, its best to do some homework first if you plan any Linux installations on older hardware. In addition to the hardware details, an understanding of Linux partitions is highly recommended.

Type of Installation

There are several options to choose from so its best to review before starting. The Workstation and Laptop options provide manual or automatic partitioning but will write a boot record (MBR). Anyone wanting to multi-boot should not select these options. The Server option will format all partitions and does not include the X Window system. So I selected the Custom option as it allowed the most flexibility (I wanted LILO booting and manual partitioning). In addition, there is an Upgrade option. The default installation uses a graphical interface. If video problems occur, you will want to choose the text mode installation.

Partitioning Linux

I learned this one the hard way. After an initial installation, I found out some recommended partitioning schemes and then the Red Hat 7.1 particulars, so decided to re-install. Red Hat's Disk Druid program is very good. The following table shows

my partitioning on the Notebook's 6GB drive:

Device	Partition	Size (MB)	Comments on storage usage
/dev/hda1	/boot	40	
/dev/hda5	/home	1,000	User data
/dev/hda6	/usr/local	1,000	Programs and data
/dev/hda7	/usr	2,400	Programs and data
/dev/hda8	/opt	400	Star Office
installation			
/dev/hda9	/	300	
/dev/hda10	/var	400	Web and SQL server data
/dev/hda11	[swap]	200	
/dev/hda12	/tmp	100	

The /usr partition is where most programs are installed. I still have 1 GB free after installing the programs I use.

Booting with Lilo

Linux provides a very good and easy to use booting facility using Lilo. On the dual OS desktop unit that had an existing Windows 2000 installation, the desktop boots to Linux and a menu appears: I can select Linux (now my default) or Windows.

Selecting Packages

The Red Hat package installation was outstanding - you basically push a button and 500 or so rpms are installed and ready to use. I choose Individual rather than Group package selection. The list of software to select from is very extensive. And when you select a package, RPM checks what other packages are needed and installs them. Since the installation, I have installed other packages using Red Hat's Package Manager (RPM) facility - one of the best features of Red Hat.

In addition to the installation options outlined above, you will need to select your language (18 included), keyboard, mouse and graphic facilities, then configure network, firewall, timezone, account and authentication (by default, all passwords are encrypted in Linux). To give you an idea of how easy this can be, the firewall security options are High, Medium or No firewall.

RED HAT CONFIGURATION

Hardware Configuration

This was one of the biggest challenges with installing Red Hat Linux on older hardware. Red Hat's Xconfigurator program frequently does not work on computers 2 - 3 years old! To get some hardware to work, manual editing of configuration files was needed. Red Hat Linux is years behind Windows in this area.

Package Configuration

The Red Hat installation facility provided packages that worked correctly the first time. However, many services required configuration, either using a GUI front end, or modifying the configuration files manually. For example: network, Samba, Apache and MySQL all needed some basic set-up work. Reasonable documentation on how to set-up these services were usually available and easy to find.

RED HAT - THE GOOD, THE BAD AND THE UGLY

THE GOOD...

There are a host of reasons why Linux is a feasible replacement for Windows 2000, but the primary feature is that the "open source" applications are free! Here is a summary of the best Linux has to offer, based on my personal experiences:

Software License Fees

Free Open Source Software. Free Red Hat downloads are available but obtaining Red Hat on CD-ROM is recommended; the downloads could exceed 400 MB of files.

Service and Support Fees

Different packages are available in USA. Prices range from US\$40 - \$180. To obtain in Thailand, add at least US\$100 for shipping and custom fees:

- Standard Workstation: 30 days support for 1 system. Includes 2 Cds and 1 manual.
- Deluxe Workstation: 60 days support for 5 systems. Includes 3 Cds and 2 manuals.
- Professional Server: 90 Days support for 10 systems. Includes 9 CDs and 4 manuals.

The best of Red Hat 7.1

Red Hat 7.1 was very stable by default, fast and configurable, using Linux 2.4.2-2 kernel. As with Windows 2000, you may crash an application, but not the OS!

- Many great GUI applications available that provide similar functionality to Windows, in particular, KDE desktop and StarOffice. But a few Windows application are missing.
- Excellent Web server, SQL server and server scripting environment (Apache, MySQL, PostgreSQL, PHP, Perl) for developing dynamic web pages.
- Red Hat's Installation Program - very polished. Plug-n-Play detection very good, but not perfect. And all the applications installed worked, once any required configuration steps were completed.
- RPM - Red Hat Package Manager. Both command line and GUI versions. When installing a package, RPM checks what other packages are needed and provides a warning.
- Lilo - Linux boot program with multi-boot options.

Linux Documentation

Some Linux documentation (from the Web or off the CDs) is excellent, especially the HOWTOs and Guides. MAN pages are sometimes quite helpful. The 4 Red Hat manuals, both printed and electronic formats, are useful. And there are many great books available.

Linux File Compatibility

Excellent. Linux can read and write Windows FAT and FAT32 files, whether stored on hard drives, network drives or floppy, ZIP and LS-120 disks. Some of the Linux image applications will read and write virtually any graphic file. Archive files (ZIP and tar.gz) are fully compatible on both OSs. And Star Office provides compatibility with Office 97/2000 files.

XFree86 4.0.3 and KDE 2.1.1

K Desktop Environment: Much more than just a window manager for XFree86, KDE is a complete GUI Desktop Environment. Similar to Windows 2000, its highly functional and configurable.

Star Office 5.2

Thanks to Sun Microsystems, Star Office is free, runs on Linux and Windows 9x/NT/2000, and reads/writes Office 97/2000 documents! So files can be easily shared in the office even if there are different OSs and/or different Office suites.

Samba

Great file and printer sharing facility for mixed Windows and Linux networks using SMB protocol. Includes SWAT: a browser-based interface for administering Samba.

Non-Red Hat Downloads

You can find many excellent applications for Red Hat on the Web. Its best to look for binary RPMs for Red Hat 7.x. I use and recommend the following:

- Opera 5.0. An excellent browser. A 1.6 MB download.
- Webmin 0.87. A comprehensive browser-based interface for administering Linux. A 4 MB download.

THE BAD...

Some of these problems would only be applicable when Linux is installed on existing office computers. Therefore, this section is divided into Installation and Configuration problems (not relevant to new computers with pre-installed Linux) and Operational issues:

Installation and Configuration Problems

Hardware Configuration

Most hardware was detected automatically. However, to get some hardware to work, additional configuration is required that frequently requires the manual editing of config files. If you know what to do, its easy. But finding out how to complete these configurations can be difficult and time consuming.

XFree86 4.0.3

Graphic configuration problems may occur during installation, especially with notebooks, and certain graphic chipsets and/or monitors. Despite repeated attempts to improve the desktop's display, XFree86's graphical display (1024x768 @ 16M colors) was not as good as Windows 2000 using the same desktop hardware. Windows uses 96dpi fonts while Linux uses 75dpi fonts. I tried every tweak and troubleshooting tip I could find, but none worked. While this is a common complaint, some users report improved displays.

Operational Issues

KDE Applications

Many of the KDE applications, including Konqueror, are great, but I found Konqueror and K Package Manager to be unstable at times. Some of the main problem or missing feature issues include:

- File Manager does not sort files correctly - folders always come first. Windows Explorer correctly sorts my files by date when requested, either in ascending or descending order.
- Tools - File Find: cannot return to Search Results, and crashes all too often.
- Buggy file selection using Ctrl key - frequently selects first directory.
- No auto-complete or history of what you entered in forms.
- No Thumbnail views.
- No "Select All" files option.
- No "Snap To" mouse pointer option (Automatically move pointer to default button in dialog boxes).
- JavaScript support poor.

Numeric Keypad

The following keys on the numeric keypad of industry standard PC keyboards do not always work as they do in Windows:

- Enter and Del keys only work sometimes
- For selecting text, Shift + RightArrow or LeftArrow, never work

Star Office 5.2

Windows 2000 or Linux KDE provide great GUI desktops. Star Office has yet another desktop with its own fonts, mouse buttons and pointer, etc. Requires a minimum 35 MB of physical memory to run - no option to run StarWriter only, even if the only package installed is StarWriter.

Red Hat Linux

Could not get CD-ROM or LS-120 drives to auto-mount. The mount/umount commands needed to be executed to mount and un-mount these drives.

Red Hat's RPM

The RPM databases get corrupted sometimes with

"segmentation fault" crashes. The rebuild database facility corrects this problem sometimes; on one system, I was forced to restore from a backup. RPMs from non-Red Hat distributions or from older versions - they frequently do not work.

The Red Hat installation was fully automated once the packages have been selected. However, after the installation, if you want to install a package that is dependant on other programs, nothing is automatic. You need to identify what packages are needed, hunt them down and install them before you can install your program. Windows is much easier. And be careful using non-Red Hat RPMs

Red Hat Support

If you do not purchase a Red Hat package and a support program, Red Hat will not answer Email questions. The only free support available is from user-supported web forums, and many of these requests never get answered.

Inconsistencies

Using Console programs, very little consistency with file naming or commands, eg.

- To get help, here are the possibilities: -help, --help, -h, -?
- Configuration files: no suffix, or one of these: cfg, cnf, conf, ini, config.
- Gzipped Tar files: tar.gz, tar.z , tar.Z, tgz

THE UGLY...

Installation and Configuration Problems

This is the dark side of attempting to install Linux on existing hardware. Not relevant to new computers that come with Linux pre-installed. Refer to the update section below for updates to these problems.

Xconfigurator

This XFree86 configuration program frequently does not work on older computers! When problems arise, you need to edit the configuration files manually. Until Linux provides an automated and reliable graphical configuration facility, most business users will not even consider Linux as a viable alternative to Windows 2000, even though Linux offers many advantages, like being able to utilize older hardware. This is the worst of Linux!

hwclock

Red Hat Linux 7.1 had failed miserably at understanding the hardware or BIOS clock, the system clock and the local timezone on the desktop unit. After several frustrating weeks, multiple Internet searches of HOWTOs, TIPS and guides, about 144 dozen reboots, it works and then it does not work! Unbelievably buggy software!

Operational Issues

These problems may not be issues with many business users. However, for programmers, Web developers or IT professionals, these can be critical areas of concern.

File timestamps

When copying files under Linux, original timestamps are replaced with the current date. So the "date last modified" file attribute becomes "date last copied". This becomes a nightmare for anyone dealing with many files - how can you keep track of when a file was last modified. You can force the original timestamps using cp -p., but this means not using the GUI file manager. Very poor Linux design feature!

Upgrading KDE

KDE 2.2 was released in late August, so to evaluate a major upgrade using binary RPMs for Red Hat 7.x, I downloaded the files from SourceForge (37 separate rpm files in 3 different locations totaling 69 MB). The downloads took 10 hours over 2 days. But after the upgrade, KDE refused to run. I spend many frustrating hours trying to recover from this upgrade failure, but in the end, I decided to scrap it, and re-installed Red Hat. Note that Red Hat did not provide the RPMs used.

And here lies one of the biggest challenges in the Linux world. The Red Hat installation is outstanding - you basically push a button and 500 or so rpms are installed and configured correctly. But upgrading individual applications, especially for a large package like KDE, is far from pushing a button. Upgrades in Linux have a long way before they will be as easy as upgrading Windows applications.

Documentation

Some of the Linux documentation is inaccurate, incomplete or missing. Many authors simply say "I do not have the time to write documentation", or "The program is self-explanatory". Give us a break!

UPDATE: RED HAT RE-INSTALLATION

After the failed KDE 2.2 upgrade, I decided to re-install Red Hat 7.1 on the Asus test system. For a background on the first installation, This report details only to configuration problems relating to XFree86 and the date/time, hich were both solved!!

XFree86 Graphical Configuration

Xconfigurator

During the installation, Red Hat 7.1 correctly detected the graphic card (Asus AGP-V3100 using S3 Savage chipset) and the monitor (15" ADi ProVista). However, Xconfigurator could only successfully test at 256

colors and did not install the SVGA driver. So the following manual procedure was used to obtain 1024x768 at 16 million colors using the SVGA driver:

- XF86Config. All devices and screens were removed except the following:
 - Device: S3 Savage3D with Option "xaa_benchmark"
 - Screen: Two drivers were defined - "svga" and "accel"
- XF86Config-4. The monitor, device and screen sections had 1024x768 modes and 16 color depths.
- The X link was set to XF86_SVGA. The original X was linked to Xwrapper.

And on a re-boot, the 1024x768 screen in 16 millions colors was available. However, after this procedure, the XFree86.0.log file was no longer produced. My guess is that if I had linked Xwrapper to the svga server, the log would be available. I could find very little documentation on Xwrapper. Several readers have suggested using other X setup programs: XF86Setup, xf86config, xf86cfg and XFree86-configure. The only X setup programs available with the Red Hat package, in addition to Xconfigurator, that I could find were:

- xf86config: this did not work.
- xf86cfg: refused to read X config files.
- XFree86-configure: this did not work.

100 dpi Fonts

Previously, I had changed the catalogue and default resolution in /etc/X11/fs/config so that 100 came first. But this did not work.

The other change needed was in the file /etc/X11/xdm/Xservers. Here is the modified line:

- :0 local /usr/X11R6/bin/X -dpi 100 -gamma 1.6

The documentation stated that to make the new 100dpi fonts available, restart xfs and then X. But this did not work. So I ran Gnome, and for the first time, a nice display using 100dpi fonts. So I re-booted and KDE had the 100dpi fonts, including some nice TrueTypes!

Font Anti-Aliasing

I initially set this to "on" using KDE's Control Center, but was unsure if it was working. I later found out that while X reports loading the rendering extension, xdpinfo reports no rendering extension.

DATE AND TIME CONFIGURATION

I had experienced a serious problem with the date/time. On every re-boot, the date was changed, independent of any changes I made to hwclock. Once this problem started, I hunted the HOWTOs and FAQs, and that is how I got involved with hwclock.

During the re-install, I told Red Hat to not store the date as UTC and set the timezone. Since then, date has been 100%. So how did the date get totally screwed up in the previous installation?

Here are my suspicions: KDE's Control Center has a System facility, and under that, a Date & Time section. After the re-install, I checked this and timezone was set to: [no selection]. My best guess is that in the previous installation, I had used this GUI configuration tool to set the timezone, and that likely caused the mess. While I had identified the date and timezone config files under /etc, who knows where or what KDE stores.

Was hwclock "Unbelievably buggy software"? It appears that this award goes to KDE!

This is the second time I have found the System facility in KDE's Control Center to be un-reliable. Several months ago, I had used KDE's Control Center - System - Boot Manager (LILO) and found that it corrupted my lilo.conf file. I had backups and was able to restore.

The lessons I have learned on this are:

- Use GUI front-end configuration tools with caution
- Backup configuration files before using any GUI tool

DESKTOP APPLICATIONS

This section provides further details of my software reviews and comparison with Windows 2000. Some of these programs were used for testing only, while others were used daily. I have grouped the mostly desktop applications by Red Hat Linux, KDE, Non-KDE, Star Office and Non-GUI.

Red Hat Linux

OS & GUI

It's difficult to compare Operating Systems when the GUI desktop in Windows 2000 is the OS, while in Linux, the OS is a robust, powerful and compact text-based operating system. Linux uses the X Window system for graphical displays, provided by XFree86.

Linux Administration Tools

Linuxconfig comes with Red Hat, but is quite buggy. There are separate GUI fronts-ends for most administration tasks, but the best I found was Webmin 0.87 (<http://www.webmin.com/>) - a very powerful browser-based interface for doing all your Linux administration. Not included with Red Hat, you will need to download (4 MB).

LAMP development

This Web development facility is based on "open source" software and is the without a question the best in Linux!

- Apache Web server. The majority of Web servers use this software.
- MySQL database server. Fast and reliable. A high-end SQL server, PostgreSQL, is also included with Red Hat.
- PHP scripting. Similar to ASP, but easier to learn, this is an excellent server scripting facility.

The above PHP server environment can be run on Windows. I downloaded PHP Triad 2.11 from Source Forge. (<http://sourceforge.net/projects/phptriad/>)

KDE Applications

KDE's GUI desktop environment and applications are excellent.

KDE 2.1.1

While there are various GUI window managers for XFree86 available in Red Hat, I selected KDE - a complete and highly customizable desktop environment with many applications. It includes file management, easily configured menus, utilities galore, and a familiar desktop feel, which is equal to or improves on Windows' look and feel.

- Konqueror (KDE's version of Windows Explorer/Internet Explorer). Some very nice features not found in Windows! And very easy to customize.
- Excellent mouse select and paste. Whatever text is selected from anywhere (except Star Office), it can be easily pasted anywhere with a single mouse click!
- KDE Control Center is very useful: it has a Windows-like 'device manager' and a Look & Feel section to configure KDE the way you want it.
- I use KWrite 2.0 for editing both text and html files; its a very good editor. And for accessing the Web, I use and like KPPP and KMail for dialing my ISP and checking/sending POP Email. I also like Konsole (the Linux Console).

My customized KDE desktop is better than Windows 2000!

Browsers (File and Web)

I like Konqueror. Its fast and very handy for file management and browsing local files. But it is not a good Internet Browser - it crashes all too frequently and provides limited JavaScript support.

Mozilla 0.7 and Netscape 4.7 are included with Red Hat, but I prefer Opera 5. This is a commercial program available as a free ad-supported version. Download is only 1.6 MB. If you pay to register (US\$40), the advertising disappears. Highly recommended browser!

Archive Utilities

WinZip 8 on Windows and Archiver on Linux are fully compatible with both ZIP and tar.gz files. These archives can be copied and used on both Windows and Linux. One feature of WinZip that I use regularly is selecting a folder and right-clicking - I get an option: "Add to folder.zip". I can do this with several steps under Linux, but the Windows / WinZip integration is far better.

Non-KDE Applications

The applications use the inferior Gnome-style GUI.

Gnome Window Manager

I have tried Gnome, but did not like it. The graphical interface is poor compared to Windows or KDE: the windows frequently do not scale, the buttons are difficult to understand and the File Manager is quite inferior when compared to KDE's Konqueror.

CD-R/RW Recorder

Adaptec's Easy CD Creator is the industry standard for burning CD-Rs in the Windows world. It's a highly intuitive program that works very well. It would be difficult to match this great program.

On Linux, I selected the GUI Xcdroast and downloaded the latest RPMs. The program detected my HP CD-Writer, and created a 600MB CD-R in record time. Not as easy-to-use or intuitive as Easy CD Creator, but it sure works!

Image Editor

GIMP 1.2 is a highly rated, robust and full-featured image-editing program for Web graphics and is frequently compared to PhotoShop. It works with all image files. I quickly made changes to files created in FireWorks 4. However, in the FireWorks file, I had 3 text layers and a bitmap layer, in addition to the background. With GIMP, the file loaded as a single layer.

Star Office Applications

Star Office provides yet another GUI desktop that I do not like - it has its own fonts, mouse buttons and pointer, etc. But it works.

Star Office 5.2

A complete office productivity suite. It's not as sophisticated as Microsoft Office 2000, but anyone looking for a complete set of free office tools for Linux, Star Office is the best choice. Fully compatible with Microsoft Word, Excel, and PowerPoint 97/2000 file formats.

No help facility on the desktop unit! The help system failed to work despite repeated re-installs, although

the help screens did appear on 2 occasions, only to disappear the next time! Possibly a font configuration issue.

StarWriter

Excellent Word Processor. I like Word 2000. There are differences between Word and StarWriter, but even without any help files, I quickly created, edited and formatted documents. This is one Linux program that business users can quickly use productivity.

Nice formatting menu available using right mouse button. Produces clean HTML files. And includes a good 'WYSIWYG' html editor. StarWriter's spelling is not as good as Windows, but its Thesaurus is excellent - for each word, it offers a descriptive phrase of what it means.

Three Word 97 files, ranging from 75 to 460 KB, were loaded into StarWriter. Formatting, including tables and graphics, worked great.

StarSchedule

Reasonable PIM. It looks much like Outlook, can be customized and includes an adequate address book. Generally impressed. But I could not import information from Outlook: my contacts or events. If you decided to use it, be prepared to manually enter your contacts, events and tasks!

My biggest disappointment with Red Hat was the missing "Outlook" application. A good PIM with integrated E-Mail. As I use Outlook daily, I am still looking for a Linux alternative.

StarMail

Industry standard E-mail programs have 3 basic folders: Inbox, Outbox and Sent Items. But not with StarMail. It has none - you have to create an Outbox (and I never did find out about the other 2 required folders). This is one program I do not need. I use Kmail, an adequate E-Mail client included with KDE.

StarImage

Very capable Image editor, although I could not find a text tool. It will write GIF, JPG, PNG, BMP and TIF formats and in addition, read PSD, PCX, EPS and PCD.

StarImpress

Two PowerPoint 97 files, both under 100 KB, were loaded into StarImpress. One file was a slide show and displayed identical to PowerPoint. The other was a fancy diploma that displayed fine, except for the different fonts.

StarCalc

Three Excel 97 files, ranging from 45 to 300 KB, were loaded into StarCalc. Formatting and calculations worked fine. This is a very powerful and compatible

SpreadSheet.

Non-GUI Applications

These applications are frequently executed using the Linux command line (console). In some cases, GUI front-ends to these programs are available.

Red Hat Package Manager

One of the real benefits of the Red Hat distribution is RPM, a powerful console program to manager program installs, upgrades and un-installs. Developed by Red Hat, RPM is used by several Linux distributors. In addition, there are GUI front-ends: both Gnome-RPM and KDE's Package Manager are good.

Archive

The tar program is ideal for backups as it not only backs up Linux filesystem folders, but will compress them using Gzip. Very efficient. And WinZip reads these archives. KDE produces a GUI backup program called Kdat, but I could not find it with this Red Hat version.

PDF Writer

Using StarWriter, I imported a complex Word 2000 file, then printed it as a PS (PostScript) file. Then using the Linux console, I entered this command:

```
ps2pdf resume.ps resume.pdf
```

The resume.pdf file displays fine, but is 100K larger than the Acrobat Writer 5.0 version produced in Windows. Not as integrated as Acrobat Writer in Windows and not as efficient, but it works!

FAXING

I used a command-line program with this syntax:

```
fax send -v number file.txt
```

And off went a fax. Files can be text or PS. You can also receive faxes with this tool. KDE's Kfax is a GUI that will view and print fax files.

This article is re-printed with permission. The originals can be found at:

<http://www.robval.com/linux/desktop/index.asp>

Interview with Ben Collins, DPL

Author: Daniel Stone <tamriel86@hotmail.com>

INTRODUCTION

I recently interviewed Ben Collins, Debian Project Leader (via email), on his thoughts on a number of issues concerning Debian. Read on for the full text of the interview; this was one of the more interesting bits:

Any other thoughts?

I'd like to see some community research on how people think we should counter current Government trends to take away the freedoms that have allowed the free software community to flourish as it has in recent years. I think we need to cover as many ideas as possible on how to get Congress/Senate (U.S. specific in that case) to see how this would hurt us more than anything. In fact, I'd love to get a group of people to visit Capital Hill and talk directly to politicians.

(Firstly, my apologies for the lack of originality on some of the questions; it's hard not to ask them though -DS).

INTERVIEW

Firstly, how did you get started with Linux and Debian, and how long have you been using it?

I started with Linux back when I wanted to learn Un*x. It was a great way to get experience, and the fact that it was free and open sort of drew me into the whole community. I've been using it a bit over 4 years now.

What else are you involved in, in the Linux/Free Software community?

- SPARC/UltraSPARC porting (including minor gcc/glibc testing).
- Linux1394 (IEEE-1394/Firewire):
<http://linux1394.sourceforge.net/>

Were you always active as a developer, or did you stay as a user for some time before thinking about becoming a developer?

I was a user for about 3 months before becoming a developer. It's kind of hard to use Debian without becoming active in it's development.

So, how did you come to be the maintainer of arguably the most important package in any distribution (glibc), and eventually become DPL?

Regrettably, obtaining libc is not what I consider a milestone for me of any sort. A long time developer, Joel Klecker (Espy), gave the package to me the day before he died.

As for DPL, I think I worked hard to earn the respect and trust of my peers, who obviously felt I was capable of the position (or they were just having a bad day :). Winning the last election was a big confidence builder for me personally, mainly because of the caliber of developers running for the position.

What made you decide to run for DPL, and what issues did you base your platform on?

I ran for DPL for the same reason that I think anyone would; because I wanted to make a difference. My main concerns were a few loose ends in our structure, namely New Maintainer. Fortunately, the New Maintainer process has smoothed out quite a bit since I became a developer, and the fires died down.

You seceded Wichert Akkerman, in a tightly-fought race with Branden Robinson, among others. What do you think are the differences between yourself and Wichert, and even the previous leaders - Ian Jackson and Bruce Perens - in terms of leadership?

Well, personally I don't like to compare myself to anyone. Everyone one of the previous DPL's had their strengths and weaknesses, and I'm no different. Probably the one thing I'll have in common with them after my term is the feeling that I could have done more, and a new respect for the position.

What has been the highlight of your DPL term so far? So far, the fact that nothing has fallen apart :) Really, I don't have any notable high points. Real life circumstances have prevented me from putting in the time that I would like to see some things done.

What do you do in The Real World, and does your job have any overlap with Debian at all?

For money, I have several trades. Currently I am working for an IT services startup based on Linux, and I also do some contract work either developing software, or doing Linux kernel work.

And to blatantly rehash a question from the Wichert interview, how do you think Debian GNU/HURD and Debian BSD are going? Have they made any real progress?

I've no idea, honestly. I'm not sure that I should be sticking my nose where I have no business sticking it :) The folks dealing directly with these projects would surely know better.

Since then, we've already had talk of Debian Win32; how's that chugging along?

Again, no idea :)

In your term, which has covered all of the woody freeze process so far, many new architectures have been added (mips, mipsel [little-endian], s390, hppa [HP PA-RISC], ia64 [Intel Itanium], and s390). Even though not all of these will release with woody, was adding new architectures a major goal for the release?

I think we do well to add new architectures to every release. Obviously the ia64 and hppa ports releasing with woody are due in large part to Hewlett Packard's dedication to the port by hiring Debian developers to do the work. I think Debian's architecture offerings are the largest available, and definitely one of our strong points.

Except for the fact that woody has slipped far behind schedule (as, however, predicted by AJ Towns, the release manager), how do you think it's going?

Well, I think it's going well considering we are implementing an entirely new release process with testing. Most developers are used to frozen being a milestone which switches them into release mode. We don't have this right now, so the light at the end of the tunnel is a little dimmer, so to speak. I think it will be a learning process until after woody hits the CD vendors, and it will obviously take an extra amount of time.

Then there's the issue that always manages to flare up - New Maintainer. It's been 5 months since this last flared up (which involved yours truly). How do you think the process is going, given the recent controversy?

Honestly, I think the process of getting New Maintainers _into_ the project is going well. I'd like to see the process of handling the rejected/stale applications handled a little quicker.

Do you think that there are any areas in NM that could be at all improved? Has it improved since you took over as DPL?

NM has improved a great deal since I became DPL, but not because of me by any means. I think because of my outspokenness about changes in NM, that some people were afraid I would dismantle the process and rebuild it, but I haven't touched it one bit other than talking to the DAM every so often.

How do you feel Debian's presence has been? Particularly a couple of months ago, it seemed that not 3 days would pass between events Debian (or at least its developers) had a major presence at. Do you feel that, in a world where RedHat's releases are trumpeted by ZDNet as "Linux 7.1", Debian is becoming increasingly noticable? Or is it just falling by the wayside to more commercial

distributions?

Oh no, Debian is not falling off anywhere. We may not be gaining "market share", but we are surely gaining users, as more and more people start to understand free software, and what Debian is about. We have our problems as a distribution and a software project, but it's nothing we can't overcome as time goes on. Debian will always be the same project, and will always be welcome by new users...it's the rest of the world that changes, while our ideals remain the same.

I've attempted to steer clear of the WTC issue, but I suppose I have to go there now. In the light of the recent terrorist acts, the US Government wants a *global* ban on all cryptography without government backdoor access. What's your view on this?

Obviously I am against it. Not only from a developer standpoint, but as an American, I am completely against any infringement on my freedoms. I plan on taking an active role in fighting these sorts of proposed laws, and I hope other developers do as well.

What is Software in the Public Interest, and what's your role in it?

SPI is the Non-profit organization that Debian operates under. It gives us a legal entity in which to present ourselves, and provides fund management. As the DPL, I am considered an advisor to the SPI board.

Virtually no-one outside of Debian development and hard-core Free Software fanaticism knows of SPI; do you think this is a problem? If so, how can it be improved?

Yes, it is a problem, and SPI is currently making strides to counter that.

What, do you think, are Debian's main problems?

To be honest, I don't know of any problems right now that aren't being handled. There are a few policy decisions that still linger, but nothing that will tear the project apart, or hinder our work.

How would you like to see Debian improve/progress in the future?

Just the way it always has, oblivious to corporate pressure. Debian does what it can to provide a stable and technically sounds operating system, and I think it needs to continue to concentrate on that, regardless of any trends to the contrary.

Any other thoughts?

I'd like to see some community research on how people think we should counter current Government

trends to take away the freedoms that have allowed the free software community to flourish as it has in recent years. I think we need to cover as many ideas as possible on how to get Congress/Senate (U.S. specific in that case) to see how this would hurt us more than anything. In fact, I'd love to get a group of people to visit Capital Hill and talk directly to politicians.

Ben Collins, thank you very much for your time.

And thank you.

This article is re-printed with permission. The originals can be found at:

<http://www.debianplanet.org/debianplanet/article.php?sid=413&mode=thread&order=0>

Old Kernels Never Die

They just keep on getting patched...

this brief interchange was snipped from the **linux-kernel** mailing list. It summarises, perhaps in a few modest lines, why many people get continuing enjoyment and oft-times mirth from community-collaborative software development projects (or, in layman's terms, they're just **kevel!**)

From: Mikulas Patocka
<mikulas@artax.karlin.mff.cuni.cz>
To: linux-kernel@vger.kernel.org
Subject: [PATCH] Linux 0.01 disk lockup
Date: Thu, 27 Sep 2001 15:34:11 +0200 (CEST)
Cc: torvalds@transmeta.com

Hi.

Linux 0.01 has a bug in disk request sorting - when interrupt happens while sorting is active, the interrupt routine won't clear do_hd - thus the disk will stay locked up forever. Function add_request also lacks memory barriers - the compiler could reorder writes to variable sorting and writes to request queue - producing race conditions. Because gcc 1.40 does not have __asm__("":::"memory"), I had to use dummy function call as a memory barrier. The following patch fixes both issues. Mikulas

```
diff -u -r linux-orig/kernel/hd.c
linux/kernel/hd.c
```

From: Linus Torvalds <torvalds@transmeta.com>
To: Mikulas Patocka
<mikulas@artax.karlin.mff.cuni.cz>
Subject: Re: [PATCH] Linux 0.01 disk lockup
Date: Thu, 27 Sep 2001 08:27:11 -0700 (PDT)

Cc: <linux-kernel@vger.kernel.org>

On Thu, 27 Sep 2001, Mikulas Patocka wrote:


> Linux 0.01 has a bug in disk request sorting -
> when interrupt happens while sorting is active,
> the interrupt routine won't clear do_hd - thus
> the disk will stay locked up forever.

Ehh.. Mikulas, do you want to be the official maintainer for the 0.01.xxx series? Note that much of the maintenance work is probably just to reproduce and make all the user-level etc infrastructure available..

Linus

Rune for Linux

Author: Avatar <avatar@firstlinux.net>



What do you get when you take Unreal Tournament, spice up the graphics a bit, put it in a mythological Norse setting and change the aspect from first person to third person? What you get is Rune! And, if it comes from Loki Games, you can bet it's for Linux too.

Recently released by Loki, and with the expansion coming soon, Rune is one of the more popular 3D titles to come to Linux in a good while. The anticipation of this release isn't quite as great as that for Dues Ex but, this title does have a pretty good following. And with the release not too far off that of the Windows version, Rune is a fairly up to date game. Following what appears to be Loki official package style, the purchaser receives a DVD Style case, a single game CD-ROM, and 28 page manual with color cover.

System requirements for Rune are:

- Linux kernel version 2.2.X
- Pentium II with 3D accelerator card
- 64 MB RAM required [128 MB recommended - 512M combined RAM plus swap required]
- Video card capable of 640x480 resolution
- XFree86 3.3.5 or newer at least 16bpp
- OSS compatible sound card
- Hard disk with at least 700MB of space
- Internet connection for Internet play

Note that massive memory requirement that blows the original requirements for UT right out of the water. Wow! 512MB! Compared to the 128MB recommended for both UT and the Windows version of Rune, you have to wonder just what happened during the port to require that much memory.

And do you need that much memory? From the official FAQ: YES. Rune is very content-rich, and (to

keep it playable) the data is all cached. In addition, changes to the 2.4 virtual memory system have upped the requirements, by effectively mapping physical RAM into the first block of swap*.

I'd like to back that up as well. Not bothering to read, I loaded it up on a lesser system [64MB RAM, 256MB swap]. And though it would indeed run, I would certainly not recommend playing it on such a system.

SINGLE ME OUT

In the single player mode, one plays the game out as Ragnar. Having just recently proven to his village that he is man enough to carry a weapon into battle, the hero has a grudge to settle. In the opening story he bests Sigard, who apparently is the most bad ass warrior in the village, in order to gain his place as a warrior. The manual goes over this a bit, but suffice to say, Sigard is pissed and isn't about to let you get away without a rematch. In order to really start the game, one must first defeat him in battle.

As the game progresses, other characters enter and interact with the player. The plot thickens with very nicely done cut scenes that pop up throughout the game. The opening scene alone is worth starting over just to see it again. Too bad these can't be viewed outside the game [e.g. Quake 3 Arena Tier scenes] cuz they are definitely cool. For the rest of the game, one dashes through the many levels of several different zones. The camera angle is typical of what one would see in Tomb Raider or similar 3rd person games. Basically, if you didn't like the view in Tomb Raider, you certainly won't like it now. In each level area, the play space is huge, making for a lengthy game. Puzzles litter each level as well, mainly being of the jumping type of challenge. On each zone, one will encounter new enemies to battle with and overcome.

The weapon selection in Rune is far different than most modern games. Instead of guns, knives, and rockets launchers, one uses hammers, swords, and axes. There are fifteen weapons in all, spread across three classes. The only armor to be had comes in a variety of shields. In a final effort to push the game into fantasy land, each weapon comes with a "special" attack, accessible when one uses their Rune power. Each use of this attack will reduce the player's overall amount of Rune power, so usage is limited.

MULTIPLAYER MODE

The multiplayer portion of Rune seems to be little more than an after thought, unlike its engine brother UT. Well, maybe like a "Hey, its here so why not use it?" There are a mere two styles of play available, Deathmatch or Team Deathmatch. As with UT, Rune includes a server browse, which makes getting connected that much easier. Mutators are also available, most from the original engine. However, if one is looking to get a little practice in before going out into the world, forget it. Rune does not include bots to fight against. Perhaps the upcoming add-on, The Halls of Valhalla, will fix this up.

WHINING LIKE A THREE LEGGED MULE

While I must be in the minority, I had some trouble just getting the game to run correctly. I'm not talking about my little mistake with the memory. I installed Rune on my VectorLinux 2.0 box, and it ran great until I tried to save. It crashed right out every time. Loading someone else's saved game resulted in the same thing.

I tried Mandrake 8.0, but the game wouldn't even load there. Someone later discovered a symlink issue that caused this, and proposed a fix for it in Loki's Fenris bug tracking system.. However, they still couldn't save games. Personally, I had to go all the way back to Mandrake 7.2 in order to get the game to run correctly.


Not good. But like I said, only one other person reported as to be having the game problem [also in Mandrake 8.0]. Your mileage may vary.


CONVINCED YET?

So, are you going to run right out and buy Rune? Unless you're really opposed to 3rd person style games, great 3D graphics, and/or an involving experience, you should. Be forewarned though, your system better be up to snuff or expect some slow game play. Being based on the UT engine, Rune accepts many of the same tweaks, which can help smooth out the game and give it even better visual quality. While not what I would call a breakthrough game, it is certainly entertaining and should keep you busy for a good while. With the expansion pack having just been released by Loki, expect to have even more fun. With luck, we'll even have an expanded review sometime soon.

This article is re-printed with permission. The originals can be found at:

<http://www.evil3d.net/reviews/software/rune/>





Testing a Router or Firewall

Author: Laurent Constantin <laurent.constantin@aql.fr>

Laurent Constantin walks us through testing a router or firewall using the networking tool lcrzo.

"We do not present a generic high level method. Indeed, every network administrator has its own way to set up a device or to solve a problem. We emphasize on the low level method by presenting key points, which can be linked together in order to construct the high level method of every person."

INTRODUCTION

The life of a router is punctuated by several important steps :

- configuration
- securisation
- problem resolution
- rules' modifications
- replacement by a new router

In each case, network tests are needed to validate administrator's choices. For example :

- verify the router is working as expected
- verify an intruder cannot reach a given IP address or a forbidden port
- solving a network problem to verify why an application isn't functioning
- verify the validity of new rules
- simulate a testing network to help configuring a new network device

The aim of this article is to present a method for testing routers and firewalls. In this document, we only use the term "router", but the described method can also be applied to stateful inspection firewalls, packet filtering firewalls, proxy firewalls, etc.

BASIC IDEA

The basic idea about tests is to send a packet on one side of the router. Then, on the other side :

- if this packet is allowed, verify its arrival
- if this packet is evil (deny, drop), verify nothing is forwarded
- if this packet is blocked (reject), verify nothing is forwarded and an ICMP error message is sent back to sender (there is a rate limiting for ICMP)

METHOD

We do not present a generic high level method. Indeed, every network administrator has its own way to setup a device or to solve a problem.

We emphasize on the low level method by presenting key points, which can be linked together in order to construct the high level method of every person. The key points are :

- How to send a packet ?
- How to verify the arrival of a packet ?
- How to simulate a testing computer ?

Using only these 3 key points, an administrator can test a network.

TOOL USED

In order to illustrate examples, the free network testing tool lcrzoex is used. Lcrzoex is available at :

- <http://www.laurentconstantin.com/us/lcrzo/> [main server]
- <http://go.to/laurentconstantin/> [backup server]
- <http://laurentconstantin.est-la.com/us/lcrzo/> [backup server]

Other tools or network libraries (tcpdump, snoop, ipssend, netcat, telnet, libnet, libpcap, etc.) can also be used, but I choose lcrzoex because it incorporates all we need in one program.

HOW TO SEND A PACKET ?

Well, it depends on protocols supported by our network. For example, in this paper, we focus on IP over Ethernet, and ICMP/UDP/TCP.

WHICH KIND OF PACKETS TO SEND ?

The kind of sent packets depends on several factors:

- Do we want to check an allowed, or a forbidden flow ?
In the first case, we send a valid packet.
In the second case, we send an invalid or forbidden packet.
- Does the router is in a testing platform or in its real place ?
In the second case, it might be dangerous to send invalid packets.
- Does the sender and destination computers are available, or need to be simulated ?
In the first case, we spoof at IP level.
In the second case, we have to spoof at Ethernet level.
- Etc.

What's the difference between IP level and Ethernet level ?

When a packet is sent at IP level, the local IP stack:

- set the source Ethernet address to the MAC address of the network board
- gets the Ethernet address of the router by sending an ARP request (or by looking in its local ARP cache)

So, the user doesn't have to bother with Ethernet. However, this method doesn't allow to do tricky stuff with Ethernet, and sometimes with IP itself because the sender IP stack rejects the packet.

When a packet is sent at Ethernet level, it is directly sent on the network without going through the IP stack. It is more complicated because the user has to set Ethernet addresses, but can allow to do tricky things. To obtain the Ethernet address of a computer on the LAN, use "lcrzoex 2":

```
# lcrzoex 2 192.168.10.2
ip address : 192.168.10.2
ethernet address : 00:40:95:46:11:23
```

Don't forget that the destination Ethernet address has to be set to :

- the Ethernet address of the router, if the destination IP address is not on the LAN
- the Ethernet address of the destination IP address otherwise

HOW TO SEND AN ICMP PACKET AT IP LEVEL ?

We can use "lcrzoex 65":

```
# lcrzoex 65
source address [255.255.255.255]: 192.168.10.1
destination address [1.2.3.4]: 192.168.11.3
IP options []:
type (between 0 and 255) [8]: 8
code (between 0 and 255) [8]: 0
```

packet's data ['hello' 0D 0A]: 12345678 'my data'

version	ihl	tos	totlen	
4	5	0	0023h=	35
id		xxBfMf		fragoffset
822Bh=33323		0_0_0		0000h= 0
ttl		protocol		header checksum
80h=128		01h= 1		235Ah
source				
192.168.10.1				
destination				
192.168.11.3				
ICMP		type	code	checksum
12 34 56 78 6D 79 20 64 61 74 61		08h= 8	00h= 0	3F01h=16129
# .4Vxmy data				

The example 66 is identical, but takes its parameters from command line :

```
# lcrzoex 66 192.168.10.1 192.168.11.3 8 0
```

How to send an ICMP packet at Ethernet level ?

We can use "lcrzoex 68" :

```
# lcrzoex 68
send on which device [eth0]:
source address [aa:bb:cc:dd:ee:ff]:
00:40:33:E0:2C:42
destination address [ff:ff:ff:ff:ff:ff]:
00:40:95:46:41:BC
source address [255.255.255.255]: 192.168.10.1
destination address [1.2.3.4]: 192.168.11.3
IP options []:
type (between 0 and 255) [8]: 8
code (between 0 and 255) [8]: 0
packet's data ['hello' 0D 0A]: 12345678 'my data'
```

ETH				
00:40:33:e0:2c:42 vers 00:40:95:46:41:BC type : 0x0800				
IP				
version	ihl	tos	totlen	
4	5	0	0023h=	35
id		xxBfMf		fragoffset
F3D2h=62418		0_0_0		0000h= 0
ttl		protocol		header checksum
80h=128		01h= 1		B1B2h
source				
192.168.10.1				
destination				
192.168.11.3				
ICMP		type	code	checksum
12 34 56 78 6D 79 20 64 61 74 61		08h= 8	00h= 0	3F01h=16129
# .4Vxmy data				

The example 69 is identical, but takes its parameters from command line :

```
# lcrzoex 69 eth0 00:40:33:E0:2C:42
00:40:95:46:41:BC 192.168.10.1 192.168.11.3 8 0
```

HOW TO SEND AN UDP PACKET AT IP LEVEL ?

We can use "lcrzoex 37" :

```
# lcrzoex 37
source address [255.255.255.255]: 192.168.10.1
destination address [1.2.3.4]: 192.168.11.3
IP options []:
source port (between 0 and 65535) [2345]: 1234
destination port (between 0 and 65535) [53]:
packet's data ['hello' 0D 0A]:
Do you want an Empty string or the Default string
? (key eEdD) [d]:
```

IP				
version	ihl	tos	totlen	
4	5	0	0023h=	35
id		xxBfMf		fragoffset
31C2h=12738		0_0_0		0000h= 0
ttl		protocol		header checksum
80h=128		11h= 17		73B3h
source				
192.168.10.1				
destination				
192.168.11.3				
UDP		source port	destination port	
68 65 6C 6C 6F 0D 0A		04D2h= 1234	0035h= 53	
		length	checksum	
		000Fh= 15	1795h= 6037	
# hello..				

The example 38 is identical, but takes its parameters from command line :

```
# lcrzoex 38 192.168.10.1 192.168.11.3 1234 53
```

HOW TO SEND AN UDP PACKET AT ETHERNET LEVEL ?

We can use "lcrzoex 40" :

```
# lcrzoex 40
send on which device [eth0]:
source address [aa:bb:cc:dd:ee:ff]:
00:40:33:E0:2C:42
destination address [ff:ff:ff:ff:ff:ff]:
00:40:95:46:41:BC
source address [255.255.255.255]: 192.168.10.1
destination address [1.2.3.4]: 192.168.11.3
IP options []:
source port (between 0 and 65535) [2345]: 1234
destination port (between 0 and 65535) [53]:
packet's data ['hello' 0D 0A]:
Do you want an Empty string or the Default string
? (key eEdD) [d]:
```

ETH				
00:40:33:e0:2c:42 vers 00:40:95:46:41:BC type : 0x0800				
IP				
version	ihl	tos	totlen	
4	5	0	0023h=	35
id		xxBfMf		fragoffset
F2CBh=62155		0_0_0		0000h= 0
ttl		protocol		header checksum
80h=128		11h= 17		B2A9h
source				
192.168.10.1				
destination				
192.168.11.3				
UDP		source port	destination port	
68 65 6C 6C 6F 0D 0A		04D2h= 1234	0035h= 53	
		length	checksum	
		000Fh= 15	1795h= 6037	
# hello..				

The example 41 is identical, but takes its parameters from command line :

```
# lcrzoex 41 eth0 00:40:33:E0:2C:42
00:40:95:46:41:BC 192.168.10.1 192.168.11.3 1234 53
```

HOW TO SEND A TCP PACKET AT IP LEVEL ?

We can use "lcrzoex 48":

```
# lcrzoex 48
source address [255.255.255.255]: 192.168.10.1
destination address [1.2.3.4]: 192.168.11.3
IP options []:
source port (between 0 and 65535) [2345]: 1234
```

```

destination port (between 0 and 65535) [53]: 80
bit syn (between 0 and 1) [0]: 1
bit ack (between 0 and 1) [0]:
bit rst (between 0 and 1) [0]:
seqnum (between 0 and 4294967295) [3145138187]:
acknum (between 0 and 4294967295) [2039479918]: 0
TCP options []:
packet's data ['hello' 0D 0A]:
Do you want an Empty string or the Default string
? (key eEdD) [d]: e

```

IP		version		ihl		tos		totlen	
		4		5		0		0028h= 40	
		id		xxDMF		fragoffset			
		344Eh=13390		0_0_0		0000h= 0			
		ttl		protocol		header checksum			
		80h=128		06h= 6		712Dh			
		source		destination					
		192.168.10.1		192.168.11.3					
TCP		source port		destination port					
		04D2h= 1234		0050h= 80					
		seq num		ack num					
		BB770000h=3145138187		0					
		dataoff		window					
		5		05DCh= 1500					
		checksum		urgent pointer					
		DC0Eh=56334		0000h= 0					

The example 49 is identical, but takes its parameters from command line :

```
# lcrzoex 49 192.168.10.1 192.168.11.3 1234 80 1 0
0 2222222 0
```

How to send a TCP packet at Ethernet level ?
We can use "lcrzoex 51" :

```

# lcrzoex 51
send on which device [eth0]:
source address [aa:bb:cc:dd:ee:ff]:
00:40:33:E0:2C:42
destination address [ff:ff:ff:ff:ff:ff]:
00:40:95:46:41:BC
source address [255.255.255.255]: 192.168.10.1
destination address [1.2.3.4]: 192.168.11.3
IP options []:
source port (between 0 and 65535) [2345]: 1234
destination port (between 0 and 65535) [53]: 80
bit syn (between 0 and 1) [0]: 1
bit ack (between 0 and 1) [0]:
bit rst (between 0 and 1) [0]:
seqnum (between 0 and 4294967295) [4293488417]:
acknum (between 0 and 4294967295) [1610530550]: 0
TCP options []:
packet's data ['hello' 0D 0A]:
Do you want an Empty string or the Default string
? (key eEdD) [d]: e

```

ETH		00:40:33:e0:2c:42 vers		00:40:95:46:41:BC		type :		0x0800	
IP		version		ihl		tos		totlen	
		4		5		0		0028h= 40	
		id		xxDMF		fragoffset			
		113Dh= 4413		0_0_0		0000h= 0			
		ttl		protocol		header checksum			
		80h=128		06h= 6		943Eh			
		source		destination					
		192.168.10.1		192.168.11.3					
TCP		source port		destination port					
		04D2h= 1234		0050h= 80					
		seq num		ack num					
		FFB96F21h=4293488417		0					
		dataoff		window					
		5		05DCh= 1500					
		checksum		urgent pointer					
		A084h=41092		0000h= 0					

The example 52 is identical, but takes its parameters from command line:

```
# lcrzoex 52 eth0 00:40:33:E0:2C:42
00:40:95:46:41:BC 192.168.10.1 192.168.11.3 1234
80 1 0 0 2222222 0
```

HOW TO VERIFY THE ARRIVAL OF A PACKET ?

To check if a packet arrives on a network, we have to use a sniffer. We might however encounter problems :

- if the sniffer is on the destination computer, everything should be fine
- if the sniffer is on the way to the destination computer (on a crossed router), everything should be fine
- if the sniffer is on the same LAN than the destination computer (or if the sniffer is on a crossed LAN):
 - if the LAN isn't switched, everything should be fine
 - else, we have to connect the sniffer on a dedicated port of the switch (or saturate the switch)
- else, there is no way to see the packet

To display packets "lcrzoex 7" can be used :

```

# lcrzoex 7
Choose the print profile
1 - header and data in synthetic aspect
2 - header (without ethernet) and data in synthetic aspect
3 - header and data in array aspect
4 - header in array aspect and data in dump
5 - header in array aspect and data in mixed
6 - header and data in hexa aspect
7 - header in hexa aspect and data in dump
8 - header in hexa aspect and data in mixed
9 - personalized profile
Choose the profile (between 1 and 9) [4]: 4

```

ETH		00:40:33:e0:c2:24 vers		00:40:95:46:14:cb		type :		0x0800	
IP		version		ihl		tos		totlen	
		4		5		0		0054h= 84	
		id		xxDMF		fragoffset			
		0052h= 82		0_0_0		0000h= 0			
		ttl		protocol		header checksum			
		40h= 64		01h= 3		E502h			
		source		destination					
		192.168.10.1		192.168.10.3					
ICMP		type		code		checksum			
		08h= 8		00h= 0		AA90h=43664			
		A1 02 00 00 7B B7 E1 3A 61 57 03 00 08 09 0A 0B		#		...		[...:aW.....	
		0C 0D 0E 0F 10 11 12 13 14 15 16 17 18 19 1A 1B		#			[.....	
		1C 1D 1E 1F 20 21 22 23 24 25 26 27 28 29 2A 2B		#		...		[...:"\$%&'()*+...	
		2C 2D 2E 2F 30 31 32 33 34 35 36 37		#		...		[.../01234567	

Various displaying methods can be selected (choice 1 to 9).

HOW TO SIMULATE A TESTING COMPUTER ?

When computer A wants to reach computer B :

- computer A sends an ARP request to every computer on the LAN (it ask "what's the Ethernet address of B")
- computer B sees this request and answer saying "the Ethernet address of B is aa:bb:cc:dd:ee:ff"
- now computer A knows the Ethernet address of B, and can send IP packets to B

So, when we simulate a computer, we have to answer to ARP requests, in order to inform other computers. For example, to simulate the presence of 192.168.10.2, and saying its Ethernet address is 12:34:56:78:90:ab, we can use :

```
# lcrzoex 131 eth0 12:34:56:78:90:ab 192.168.10.2
```

FIRST EXAMPLE

In this example, we want to verify that :

- the sender 192.168.10.1 can go through

- the router (192.168.10.254 and 192.168.11.254) to reach
- the tcp port 80 of 192.168.11.3

Both 192.168.10.1 and 192.168.11.3 are "real" computers (we do not have to simulate them)

So, the testing procedure is :

- on 192.168.11.3 : sniff with "lcrzoex 7"
- on 192.168.10.1 : send a SYN with "lcrzoex 49 192.168.10.1 192.168.11.3 1234 80 1 0 0 2222222 0"
- on 192.168.11.3 : we should see the packet

SECOND EXAMPLE

In this example, we want to verify that :

- the sender 192.168.10.1 can go through
- the router (192.168.10.254 and 192.168.11.254) to reach
- the tcp port 80 of 192.168.11.4

The computer 192.168.10.1 is "real" and 192.168.11.4 is simulated with Ethernet address 12:34:12:34:12:34.

So, the testing procedure is :

- on 192.168.11.3 : sniff with "lcrzoex 7"
- on 192.168.11.3 : simulate 192.168.11.4 with "lcrzoex 131 eth0 12:34:12:34:12:34 192.168.11.4"
- on 192.168.10.1 : send a SYN with "lcrzoex 49 192.168.10.1 192.168.11.4 1234 80 1 0 0 2222222 0"
- on 192.168.11.3 : we should see the packet for 192.168.11.4

THIRD EXAMPLE

In this example, we want to verify that :

- the sender 192.168.10.2 can go through
- the router (192.168.10.254 and 192.168.11.254) to reach
- the tcp port 80 of 192.168.11.3

The computer 192.168.10.2 is simulated with Ethernet address 12:34:12:34:12:34. The computer 192.168.11.3 is real.

So, the testing procedure is :

- on 192.168.10.1 : obtain the Ethernet address of 192.168.10.254 with "lcrzoex 2 192.168.10.254" (for example, we obtain aa:aa:aa:aa:aa:aa)
- on 192.168.11.3 : sniff with "lcrzoex 7"
- on 192.168.10.1 : send a SYN with "lcrzoex 52 eth0 12:34:12:34:12:34 aa:aa:aa:aa:aa:aa 192.168.10.2 192.168.11.3 1234 80 1 0 0 2222222 0"
- on 192.168.11.3 : we should see the packet

OTHER EXAMPLES

Several other examples could be written using the same methods. This is left as an exercise for the reader.

CONCLUSION

Validating the configuration of a router is a long task, mainly if we want to verify the security of the device. The knowledge of key points, and the usage of generic tools can however simplify administrators' job. This paper only described the tests which can be done. Using these tests and they knowledge, administrators can elaborate their own method to configure or secure a router or a firewall.

This article is re-printed with permission. The originals can be found at:

<http://www.rootprompt.org/article.php3?article=2317>

Caldera OpenLinux Workstation 3.1

Author: Daniel Christle <dan@thedukeofurl.org>

Caldera traditionally has focused on the corporate market by offering a version of Linux that is aimed squarely at the corporate desktop. Their product line includes a server edition and as a result of recently acquiring SCO, a full UNIX solution.

Caldera has always been associated with ease of use having been one of the early adapters of the KDE desktop environment. They were also the first to offer a graphical install process (complete with entertainment) that was comparable with the Microsoft Windows installation. Subsequently, they became the standard which all graphical install routines in the world of Linux were measured against.

Recently Caldera have been mired in controversy, with plans to switch from open source to a more common proprietary business model. They now have per seat licencing for their distribution, the antithesis of licensing models offered by almost every other distributor of Linux based operating systems. A stance that many postulate may hurt them in the long haul. After all, if Red Hat can pull off the new business model there isn't a reason for anyone can't.

This latest release is built around KDE 2.1, and as such, contains a good many KDE development tools and the accompanying documentation. Some of the benefits being touted by Caldera include: software integration, default configurations, self hosting, secure software, system testing, and even OEM testing.

Essentially this means that Caldera has tested each piece of software included in their distribution to make sure there are no software conflicts. Every piece has been tested for proper functionality and that any OEM that bundles OpenLinux has been tested for hardware compatibility. The benefit of default configurations are the fact that Caldera has predetermined a lot of the settings for each daemon that is included in this distribution. This is so you don't have to spend hours configuring a daemon from scratch.

The stated benefit of self hosting is perhaps the most confusing for users to understand. In Caldera's bundled documentation, they define self hosting to be the "building of delivered binaries on the same system it is delivered on." This means that the source and the binaries should match and that the binaries for a Caldera system can easily be reproduced. I'm not sure that Caldera is a leg up on everyone here, as any system compiled from source code shares this benefit.

In any case, it is time to see how Caldera matches up.

SPECIFICATIONS Caldera OpenLinux Workstation 3.1

Kernel
2.4.2

C Library
GLIBC 2.2.1

GUI
XFree86 4.0.3

Compatibility
KDE 2.x, KDE 1.x, i386

File Systems
ext2, vfat, FAT, ISO9660, ReiserFS, Many more

Requirements
GUI: Pentium, 64 MB RAM min, 1 GB HD
Console: i386, 8 MB, 500 MB HD

Platform
x86

Features
Graphical Installation, KDE-Integrated Configuration Tools, LSB 1.0 Compliant

Bundled Software
KDE 2.1, OpenSSH 2.5.2, JDK 1.3, Netscape Communicator 4.77, Adobe Acrobat Reader 4.0, Quanta Plus 2.0, PERL 5.6.0, GCC 2.95.2, Many More. Installation CD-ROM

Price
Free for Download

Caldera OpenLinux Workstation 3.1 -
\$59.99 USD

TEST CONFIGURATION Test System

Processor
AMD Duron 700 Mhz

Motherboard
GigaByte GA-7ZXR Rev. 2.2

Video
NVIDIA RIVA TNT2 -AGP

Memory
256 MB Generic PC133 SDRAM

Hard Disk
Fujitsu 20 GB - 7200 RPM, ATA-100

Other Storage
AOpen 50X CD-ROM - ATA-33
Iomega Zip 100 - ATAPI
1.44 MB 3.5" Teac Floppy

Communications

AOpen FM56-ITU/2 - V.90 56K Internal ISA
D-Link 10/100 Ethernet

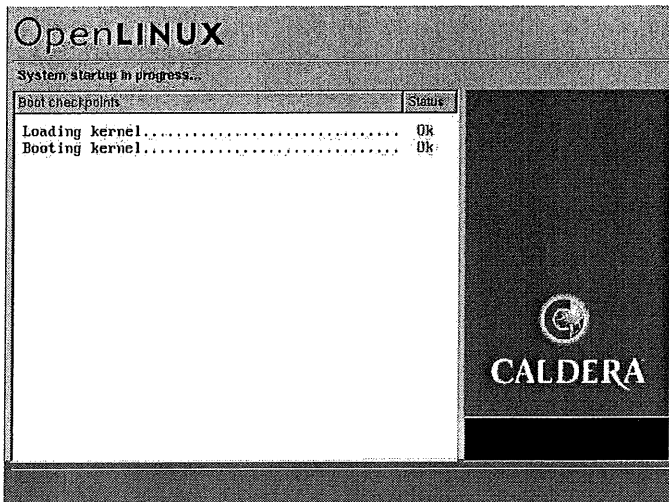
INSTALLATION

Installing Caldera OpenLinux Workstation turned out to be very interesting to say the least. For whatever reason, it did not want to install on my system without a fight. It took several attempts before it would install successfully.

Every time the install started out fine, but it would then incorrectly identified the mouse. Any attempt at changing the setting to the correct mouse definition, or any other definition for that matter, resulted in a installation failure and complete loss of control of the mouse. It kept on insisting my Logitech Cordless MouseMan was a Microsoft Intellimouse. In the end I left it at that setting just so I could complete the install. The other hiccup I encountered was the installation would fail at exactly 53%. I duplicated this six times and finally resorted to the VESA installation method. That worked and I was able to use my NVIDIA TNT2 card (a card most other distros didn't have problems with at all) normally after the install was completed by hand editing the configuration file. The text-based install also worked well in this situation.

AND SO IT BEGINS...

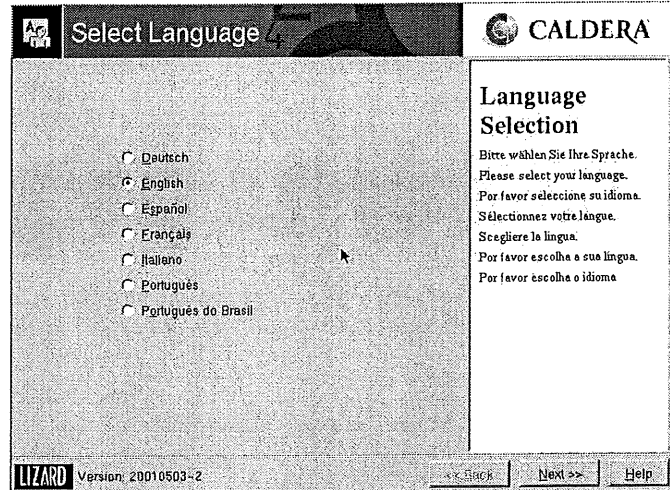
The most interesting thing about Caldera's install routine is that is graphical, and compared to other distributions, utterly amazing in look and feel. Despite the troubles I went through to get Caldera to install, the install routine was still impressive to look at.



LANGUAGE SELECTION

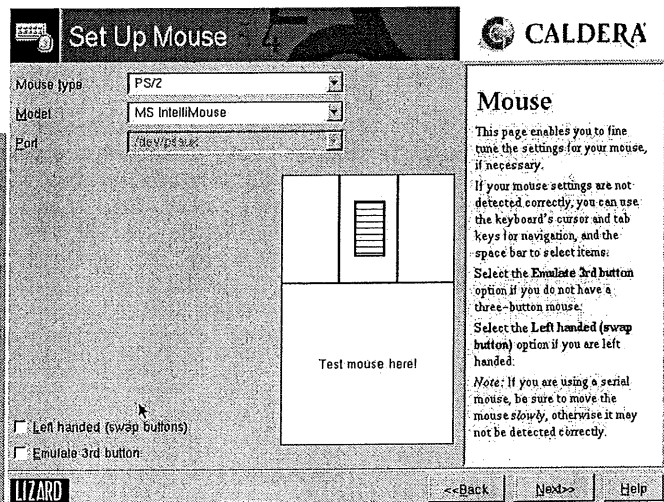
Caldera OpenLinux offers a fairly limited selection of languages when compared to other distributions, such as SuSE and Red Hat. However, it does provide

choices for English, French, Spanish, Italian, and Portuguese



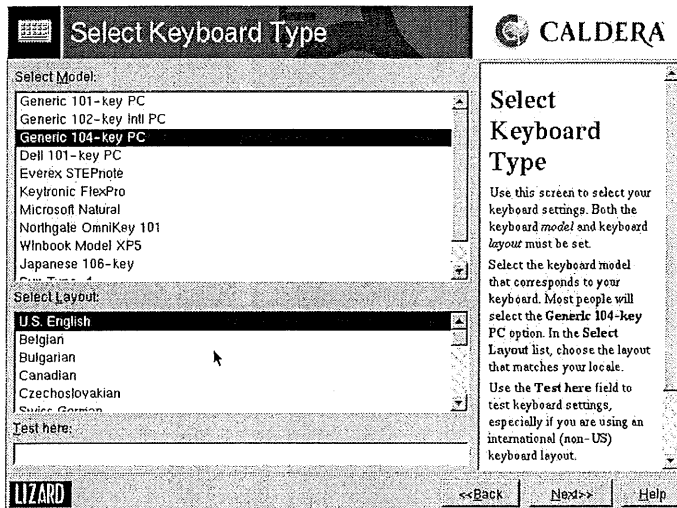
THE MOUSE

This is where things started to go wrong. Once you accept what it tells you, or can get it to accept what you tell it, you can test your mouse and move on with the installation. Incidentally, I did try Caldera on an older system with a generic serial mouse and that worked fine. Perhaps their hardware detection is out of date.



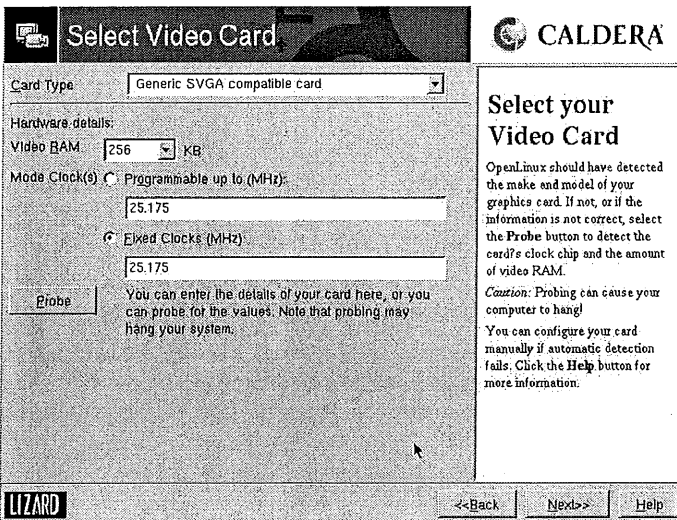
THE KEYS

Next up is the standard dialog that allows you to select your keyboard. It doesn't detect a cordless keyboard any better than it detects a cordless mouse. However, you can safely make the correct selection and everything should work well.



VIDEO

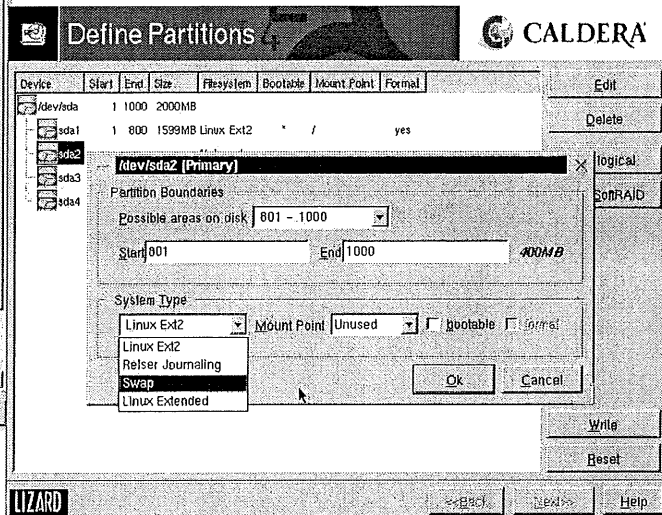
The next stage of the installation is the detection and configuration of your video card. You can select from a list in case it doesn't detect your video card correctly. A simple probe should also automatically set up the card for you.



TARGET DRIVE

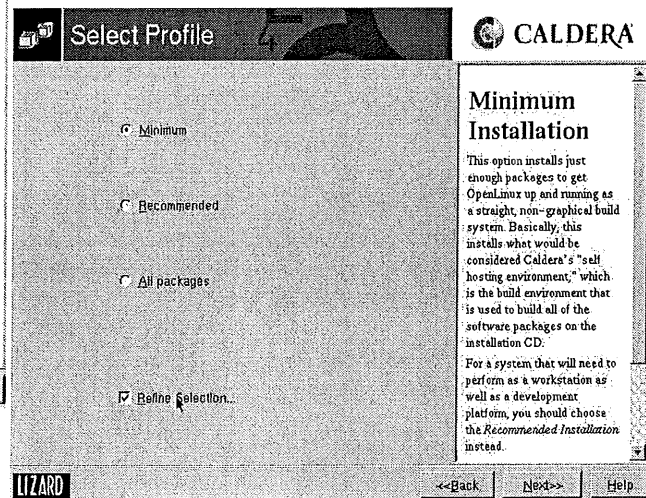
Setting up the installation target (hard drive) is a very simple matter. You can use the free space left on a drive, the whole drive, or create your own custom partitions. I chose the option for creating custom partitions. It truly is a point and click affair. It works rather well, much like the partitioning tool found in SuSE 7.2. You can choose between the ext2 (traditional) filesystem or ReiserFS (journaling). Once you have created the partitions, make sure you click the "Write" button in the bottom right corner of the screen, otherwise you won't have working partitions. Right after that, the system will ask you to format the

drives. Click the "Format" button and you will be off to the package selection portion of the installation.



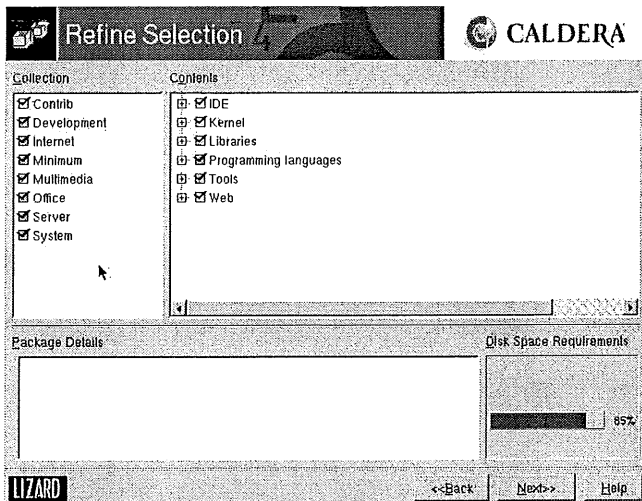
THE SOFTWARE

You have several choices when it comes to the package selection process. You can install a Minimal System (1GB), a Recommended System (1.65GB), or All Packages. If you want more control over you package selection you can check the check box at the bottom of the list, which is labeled "Refine Selection".



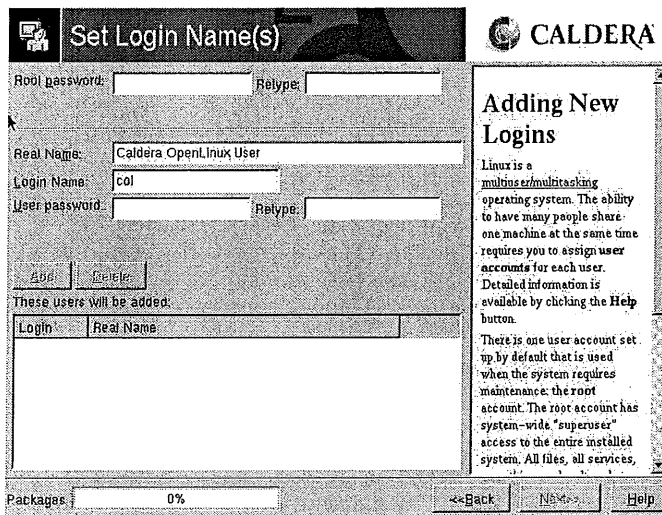
REFINEMENT

If you choose to refine the selection of software being installed you will be presented with a dialog that allows you to add and remove packages from the installation list. The list is presented to you in a hierarchical format that groups software by function.



ROOT AND USERS

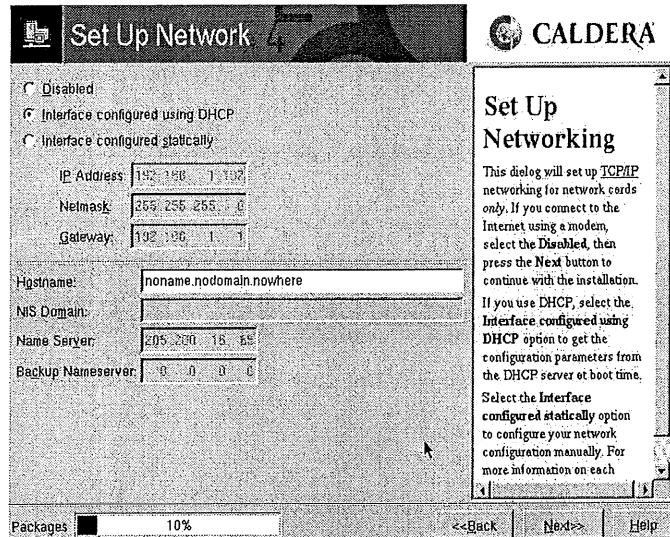
The next step in the installation process is to set the root password and create some users. You have to, at the very least, create one user before the system will continue with the installation. Notice that there is a progress bar down at the left hand side of the screen. It tells how much of the installation has been completed as you work on the latter parts of the installation process. This demonstrates Caldera's philosophy of using a multi-tasked installation routine in order to speed up the installation process.



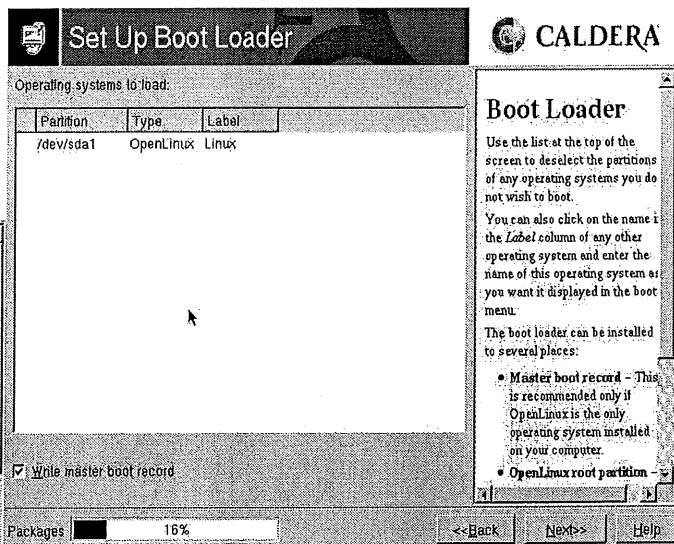
THE NETWORK

Next up is the option to set up your network connection. The network interface can be configured manually or by using DHCP. DHCP worked flawlessly on my system.

BOOT LOADER



Following the network configuration is the Boot Loader configuration. Here you can tell the boot loader what operating systems to load at boot time. If you dual boot you can choose between Linux and Windows. By default OpenLinux uses the GRUB boot manager.



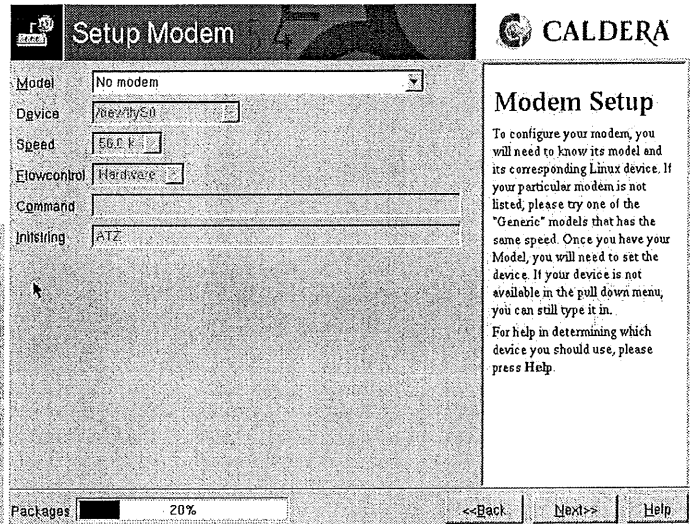
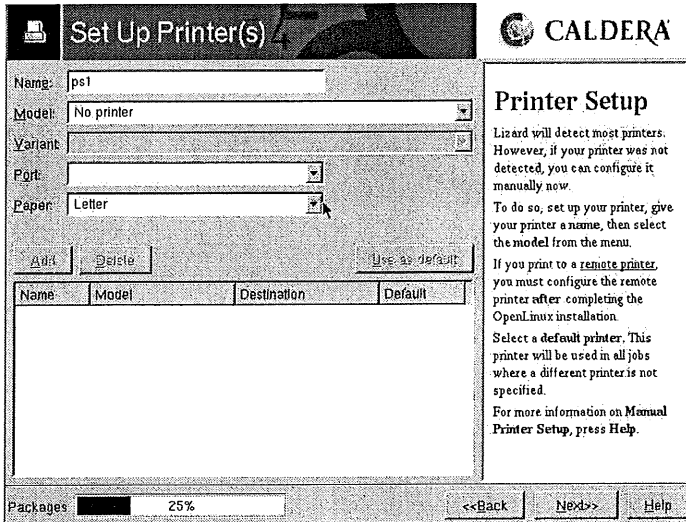
MODEM SETUP

Modem configuration is the next step in the installation process. You just need to select your port and if possible a modem definition from the list. If your modem is not listed and it is not a winmodem, chances are you can use one of the generic modems from the list. New users who are not familiar with Linux port assignments will want to use the help system to determine the correct port settings for their modem.

PRINTERS

Caldera uses CUPS (Common UNIX Printing System)

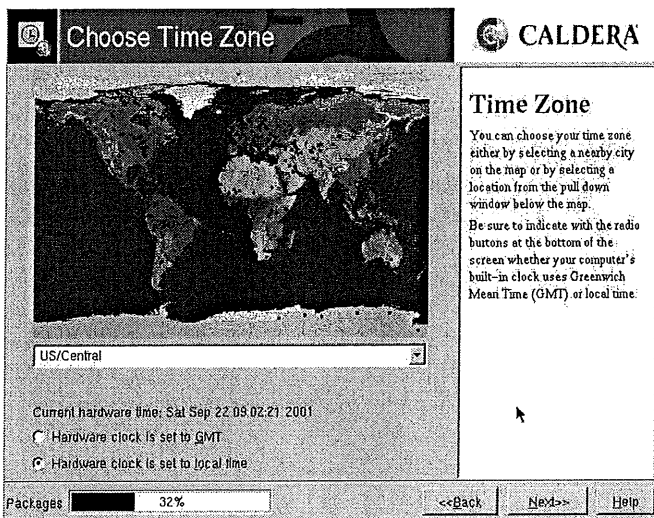
to configure its printers. Remote printers must be configured after the installation using the kCUPS interface. A big plus with CUPS is the support for a wide variety of printers. Most of the printers popular with the corporate crowd are supported here.



with Caldera OpenLinux with no problems at all. I just had to manually set up my video card.

THE TIME

Setting up the date and time for your system is very easy. In fact it appears to be the very same time zone applet Red Hat uses in their install. You can click directly on the map to select the appropriate time zone or, more accurately, select from the drop down list at the bottom of the map.



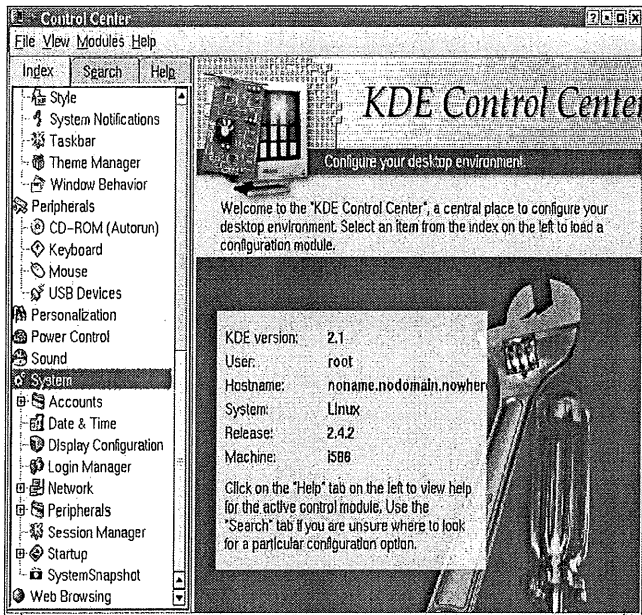
CARDS ANYONE?

After all that hard work you can play a quiet game of solitaire. Even the boss can't argue with that. Once the install is complete you need to create a boot floppy. After that the system boots and you are up and working. Interestingly enough, after all the problems I had with the install I was able to work

POST-INSTALLATION

COAS, I KNEW YOU WELL...

Anyone who has used Caldera OpenLinux before will remember the COAS administration tool that was often raved about in the past for its incredible ease of use. Caldera has now integrated the administration tools directly into the KDE Control Center. This makes much more sense than having several configuration tools scattered across your hard drive. If you are going to have graphical tools you might as well integrate them. It is much easier on the user/administrator.

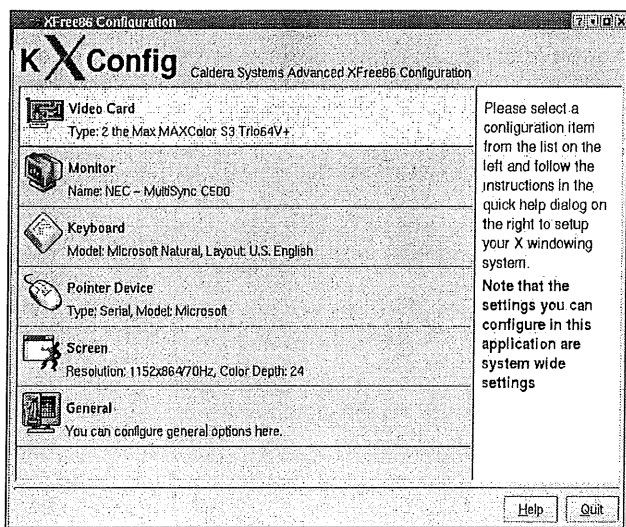


USER AND GROUP ADMINISTRATION

Setting up users and groups is very easy. Caldera does a good job here, making it much easier than the older tools that used to ship with Caldera or even the KDE User Manager (which Caldera does not ship with the distribution).

X

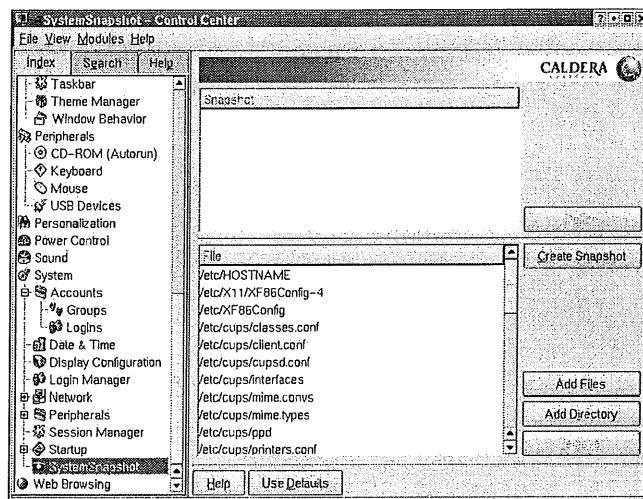
Caldera is probably one of the few distributions shipping a graphical configuration tool for XFree86 4.x. Here you can configure everything from the video card to you monitor and screen resolution. Again, Caldera makes it very much point and click affair.

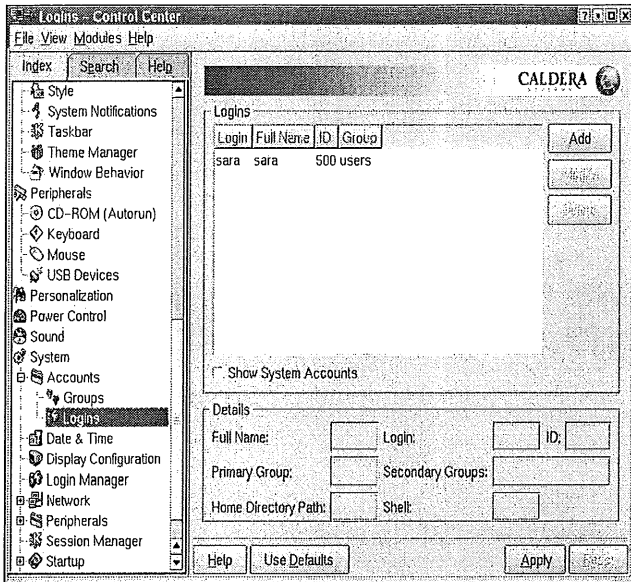


SNAPSHOTS

One of the more interesting system tools included

with this release is System Snapshot. It gives you the ability to create a snapshot of any file on your system. This allows you to very easily recover from a system failure. By being able to restore system configuration files (the default setup) you can tremendously reduce the amount of time it can take to get a system back up and functioning properly. You can also create and store a series of snapshots as you change and modify your system. Additionally, you have the ability to delete older snapshots in case they have become obsolete by your more recent changes.





OTHER OPTIONS

There is a fair amount you can do to with the configuration applets in the Control Center. You can manage the kernel modules, loading and unloading them as required. You can even set their parameters if necessary. This can be a real time saver over loading modules by hand. Other options include managing the boot loader, creating a rescue disk, managing the services that are loaded at startup, configuring sendmail, configuring dial-up networking, setting up your network interfaces, setting up DNS servers, setting up of remote and local printers, and interestingly enough, configuring a fax service.

Caldera does an excellent job with the system administration tools. Regular users are properly locked out of these tools, even. The one thing I did like about the Caldera configuration tools is that they don't overtake the system. If you need to configure a file by hand, it won't overwrite your changes. This is a real time saver and something that may be of interest to that beginning system administrator.

SECURITY

There isn't a whole lot to say about security with Caldera OpenLinux. If you want you can install some optional firewall rules. About the only comment Caldera has provided about security in the media kit they provided me was that they implement security "by closing the system by not invoking processes that invite intruders and by bundling security packages." This pretty much begs the question: why sendmail? I can only reason that it was included because of its strong corporate presence. Sendmail is a good product but there are much more secure alternatives available Caldera is certainly right about not loading too many processes. You won't find a telnet or FTP

daemon running on your system. If you need to telnet somewhere you will have to resort to SSH. This is a good thing. OpenSSH is also included both in client and server form for those that need it. About the only overt security package I could see was the personal firewall which allowed you to point and click your way through the rules setup. Beyond that Caldera doesn't appear to push security as much as some other distributions, most notably Red Hat, Mandrake (more recently), and SuSE. This isn't to say that Caldera is insecure, they just aren't pushy about it.

PROS ...

Reliability

For the most part, Caldera provides a solid and reliable Linux experience. There is no doubt that once you get past any installation woes you will have a system that works -- and works well. This is evident in their conservative choice of packages. You won't find much in the way of bleeding edge software, which has proved to burn some companies from time to time.

KDE Development

Caldera has included just about all the KDE development tools one could wish for. You get a fairly current releases of KDE Studio, KDevelop, and qtDesigner. There are other development tools available as well. Most of those are common to most distributions and Caldera has done a decent job in bring the developer the most up-to-date versions available.

AND CONS

Hardware Detection

Based on my experience with Caldera's installer I can't exactly say that their hardware detection is flawless. Compared to other distributions there were too many problems for me to give them a decent rating here.

KDE-Centric

Caldera is way too KDE-centric. You won't find another window manager anywhere and you certainly won't find GNOME on the system. Getting GNOME onto a Caldera system is quite an interesting experience, one that most users will want to undertake. Any one notice that Ximian doesn't support Caldera? Wonder why?

CONCLUSION

Usability 9.0

If anything, Caldera's OpenLinux is very usable. This is likely due to their focus on the corporate desktop where simplicity is almost a virtue.

Performance 8.0

Performance is just above par, mainly due to the lack of large amounts of software being started at boot

time. It doesn't scream like Debian or Slackware but is faster than some of the other distributions.

Innovation 6.0

I really can't say there is a whole lot of innovation happening here. This is more of a point release and anything new is really just an upgrade to what was available in previous versions.

Compatibility 7.0

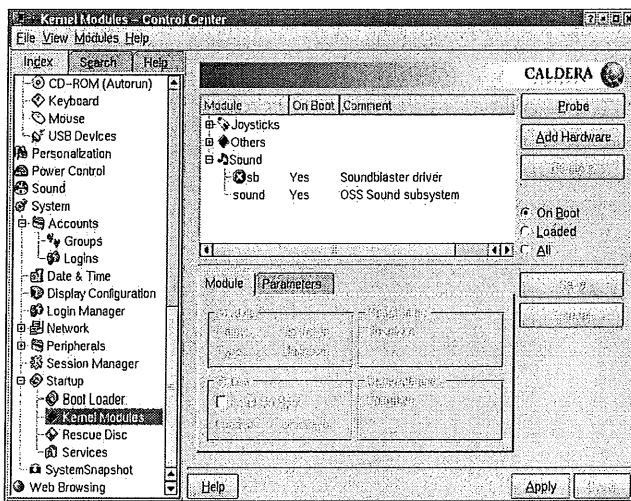
I am knocking them down a bit on this score. Their hardware detection is poor and seems to require some herculean efforts to get it working with hardware that is known to be supported by Linux. In the past Caldera has been a joy this way.

Currency 7.5

They have the latest version of KDE and assorted development tools included in this release. Where they are lacking is in the choice to ignore other desktop environments and window managers. Choice is a key strength of Linux.

Overall 7.5

Caldera OpenLinux is probably a good choice for the corporate desktop. It is simplicity which makes it very easy to administer. The exclusion of GNOME and other window managers is going to deter a lot of users. The fact that it can be hard to get GNOME working on OpenLinux at all will be a strong deterrent in itself. Companies looking to implement a Linux solution will get a solid and reliable desktop environment. However, you may want to re-evaluate the total cost of ownership as Caldera moves ahead with its per seat licensing scheme.



This article is re-printed with permission. The originals can be found at:

<http://www.thedukeofurl.org/reviews/misc/caldera31/printable.shtml>

Wasabi Systems ships update to NetBSD

Press Release

New York, NY, September 24, 2001 - Wasabi Systems, Inc., the premier provider of the NetBSD embeddable open source operating system, today announced shipment of NetBSD v. 1.5.2. NetBSD 1.5.2 includes many new and enhanced features over the previous versions. NetBSD 1.5.2, a fully functional UNIX-like operating system which runs on 44 different platforms, includes new features, such as:

- Additional machine support for new Apple iBook and PowerBook laptops.
- Security fixes for Kerberos, BIND, ssh, ntpd, ftpd, telnetd, and IP filter.
- Performance enhancements for NFS, LFS, Symbios/NCR SCSI, sendmail, and dhcpcd.
- Additional device support for Cisco Aironet; 3Com 3c555 and 3c556; Yamaha, ESS Maestro, NeoMagic 256, and CrystalClear Audio; ATA Ultra-DMA/100; and 802.1Q virtual LANs.
- Support for running Linux Vmware on NetBSD/i386.

About Wasabi Systems

Wasabi Systems, Inc. (www.wasabisystems.com), is the premier source for commercial NetBSD development, support and customization. Wasabi offers a range of integrated NetBSD system solutions focusing on the server appliance and embedded systems markets. Wasabi's team includes the world's foremost NetBSD developers, including several members of NetBSD Core Team and release engineers for the NetBSD Project.

About NetBSD

The NetBSD Operating System is an open source BSD Unix operating system descended from the Berkeley Networking Release 2 (Net/2), 4.4BSD-Lite, and 4.4BSD-Lite2 sources. NetBSD runs on 44 different system architectures featuring twelve distinct families of CPUs, and is being ported to more.

DVD players for Linux

Author: Glenn Mullikin <glmull@machineofthemoth.org>

INTRODUCTION

I used a 1.2 Ghz Athlon with 1 GB of RAM with an inexpensive DVD drive that cost me maybe \$30. So yes, MPEG2 decoding was going to be performed in software but with a machine like this, that would not be a problem, and it wasn't a problem with all of the software I tested. The quality probably doesn't equal a

standalone dvd player because there are sometimes some artifacts and sometimes the screen blanks out with one of the players and the monitor seems to think it lost a video signal but on the whole, the quality of video is acceptable. Sometimes the screen will freeze with one of the players and stay that way for a few seconds and then all of a sudden the video and audio start playing again. No, it's not perfect but we're getting close.

The real question is "Do I miss my DVD player?" The answer to that question is "I love my linux box and the more things I can do on it, the more I love it." So sure I miss Dolby Digital 5.1 surround sound on a home theater system setup but I get more satisfaction out of watching a movie on my computer than a big home theater system. Why? Because it's open source. How? Because I compiled and installed it.

SETTING THE REGION CODE OF THE DVD-ROM DRIVE

You may need to set the region code of your dvd-rom drive to whatever region of the world your dvd discs are for. I set mine to region 1 using the `regionset` command line executable that comes in the **`dvd_disc_20000215.tar.gz`** package. you can download that tar.gz from my website: http://www.machineofthethmonth.org/articles/a53/dvd_disc_20000215.tar.gz

Or it can be gotten from:

http://freshmeat.net/projects/dvd_disc/ if you want the latest version.

You have to "make" and that will produce an executable file called `regionset` that you can run. Prompts will guide you through the process when you run the `regionset` command. But beware, you can only set the region code 5 times. After that, you're stuck with whatever you last set it to, from what I hear. For most people that shouldn't be a problem since they use discs from only one region.

If you live in the United States and purchase your dvd movies here in the US, then you will probably want to set your dvd-rom drive to region 1, like I did.

THE UDF FILESYSTEM STRUCTURE OF DVDs AND HOW TO OPERATE WITH IT

Many of you are familiar with compiling the linux kernel. It turns out that DVD movies on a dvd disc are stored in a format called UDF - Universal Disk Format. <http://www.trylinux.com/projects/udf/> gives a little more information on UDF but basically, what one needs to do is compile UDF filesystem support into the kernel, either as a module or directly. I chose to compile it modularly so before I can begin working with DVD-ROM discs I need to load the `udf.o` module as follows:

```
insmod udf
```

Once I have that done, I can then put a DVD-ROM

disc (DVD movie, in this case) in my DVD-ROM drive and mount it, using the familiar tools that we all use with CD-ROM drives. From this point on, I can examine the contents of the disc as I would with a normal CD-ROM drive. Basically, from the end user's perspective, the fact that this is a UDF filesystem is a transparent thing because once the disc is mounted, it is available to look at using the standard commands. (`ls`, `cd`, etc).

So if I put a DVD disc into my DVD-ROM drive and want to mount it, I need to type the following:

```
d3:~/1scripts# mount /mnt/hdc
```

And here is what my system says is mounted:

```
d3:~/1scripts# mount
/dev/hda1 on / type ext2 (rw,errors=remount-ro,errors=remount-ro)
proc on /proc type proc (rw)
/dev/hda3 on /usr type reiserfs (rw)
/dev/hdc on /mnt/hdc type udf (ro,noexec,nosuid,nodev)
d3:~/1scripts#
```

The relevant line, of course, is the one in bold. It shows that we have a disc mounted as a `udf` filesystem type on the `/mnt/hdc` mount point. Do a man `fstab` if you don't understand how to set something like that up. The `/etc/fstab` file is where you establish links between mount points, such as `/mnt/hdc` and actual physical block devices, such as `/dev/hdc`. The fact that the device at `/dev/hdc` is a dvd-rom drive and not just a cdrom drive is of no consequence and it is not treated any differently as far as setting it up goes.

THE STRUCTURE OF A DVD MOVIE DISC

All DVD movies appear to have a similar type of directory structure, which makes it easy for you and me to mount the disc and take a look and see exactly which files are the ones that contain the video footage. These files live in the directory called `VIDEO_TS` and have the `.VOB` file extension. They also are big files, typically in the range of 1 gigabyte. Usually there are maybe 3 or 4 such files and they might be called `VTS_01_1.VOB`, `VTS_01_2.VOB`, `VTS_01_3.VOB` and so on. But whatever they are called, their naming seems to follow `VTS_0Y_X.VOB`, at least on the discs I have mounted and looked at, where `Y` and `X` are single digit numbers.

Is it really necessary to understand this particular aspect of filesystem structure on DVD movie discs? Well, it might be, just depending upon the software that you are using on Linux. With `MPlayer`, you do need to know what the `VOB` filenames are because you have to specify those. With `Videolan`, you don't. It can apparently figure out which `VOB` files are the ones for the movie and it can play them in the appropriate order, at least it did for me on the DVDs I tested so far.

But for the fun of it, let's take a look at a few examples. For example, the disc that I just mounted above in the last section, we'll look at it real quickly.

Since we already have the dvd disc mounted as a `udf`

filesystem, we can browse its contents (to see what VOB files are there).

```
d3:~/scripts# cd /mnt/hdc
d3:/mnt/hdc# ls
AUDIO_TS VIDEO_TS
d3:/mnt/hdc# cd VIDEO_TS/
d3:/mnt/hdc/VIDEO_TS# ls -l
total 4295340
-r--r--r-- 1 4294967295 4294967295 20480 May 12 1999 VIDEO_TS_0.BUP
-r--r--r-- 1 4294967295 4294967295 20480 May 12 1999 VIDEO_TS_1.VOB
-r--r--r-- 1 4294967295 4294967295 217088 May 12 1999 VIDEO_TS.VOB
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_01_0.BUP
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_01_0.IFO
-r--r--r-- 1 4294967295 4294967295 8192 May 12 1999 VTS_01_0.VOB
-r--r--r-- 1 4294967295 4294967295 28964864 May 12 1999 VTS_01_1.VOB
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_02_0.BUP
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_02_0.IFO
-r--r--r-- 1 4294967295 4294967295 8192 May 12 1999 VTS_02_0.VOB
-r--r--r-- 1 4294967295 4294967295 311296 May 12 1999 VTS_02_1.VOB
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_02_2.BUP
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_02_2.IFO
-r--r--r-- 1 4294967295 4294967295 8192 May 12 1999 VTS_02_2.VOB
-r--r--r-- 1 4294967295 4294967295 311296 May 12 1999 VTS_02_3.VOB
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_03_0.BUP
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_03_0.IFO
-r--r--r-- 1 4294967295 4294967295 8192 May 12 1999 VTS_03_0.VOB
-r--r--r-- 1 4294967295 4294967295 85338112 May 12 1999 VTS_03_1.VOB
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_04_0.BUP
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_04_0.IFO
-r--r--r-- 1 4294967295 4294967295 8192 May 12 1999 VTS_04_0.VOB
-r--r--r-- 1 4294967295 4294967295 845824 May 12 1999 VTS_04_1.VOB
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_05_0.BUP
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_05_0.IFO
-r--r--r-- 1 4294967295 4294967295 8192 May 12 1999 VTS_05_0.VOB
-r--r--r-- 1 4294967295 4294967295 196608 May 12 1999 VTS_05_1.VOB
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_06_0.BUP
-r--r--r-- 1 4294967295 4294967295 18432 May 12 1999 VTS_06_0.IFO
-r--r--r-- 1 4294967295 4294967295 8192 May 12 1999 VTS_06_0.VOB
-r--r--r-- 1 4294967295 4294967295 364544 May 12 1999 VTS_06_1.VOB
-r--r--r-- 1 4294967295 4294967295 49152 May 12 1999 VTS_07_0.BUP
-r--r--r-- 1 4294967295 4294967295 49152 May 12 1999 VTS_07_0.IFO
-r--r--r-- 1 4294967295 4294967295 8192 May 12 1999 VTS_07_0.VOB
-r--r--r-- 1 4294967295 4294967295 1073709056 May 12 1999 VTS_07_1.VOB
-r--r--r-- 1 4294967295 4294967295 1073709056 May 12 1999 VTS_07_2.VOB
-r--r--r-- 1 4294967295 4294967295 1073709056 May 12 1999 VTS_07_3.VOB
-r--r--r-- 1 4294967295 4294967295 106064864 May 12 1999 VTS_07_4.VOB
d3:/mnt/hdc/VIDEO_TS#
```

```
-r--r--r-- 1 4294967295 4294967295 26624 Aug 25 1999 VTS_01_0.IFO
-r--r--r-- 1 4294967295 4294967295 2111488 Aug 25 1999 VTS_01_0.VOB
-r--r--r-- 1 4294967295 4294967295 9646080 Aug 25 1999 VTS_01_1.VOB
-r--r--r-- 1 4294967295 4294967295 18432 Aug 25 1999 VTS_02_0.BUP
-r--r--r-- 1 4294967295 4294967295 18432 Aug 25 1999 VTS_02_0.IFO
-r--r--r-- 1 4294967295 4294967295 159744 Aug 25 1999 VTS_02_0.VOB
-r--r--r-- 1 4294967295 4294967295 32198656 Aug 25 1999 VTS_02_1.VOB
-r--r--r-- 1 4294967295 4294967295 102400 Aug 25 1999 VTS_03_0.BUP
-r--r--r-- 1 4294967295 4294967295 102400 Aug 25 1999 VTS_03_0.IFO
-r--r--r-- 1 4294967295 4294967295 159744 Aug 25 1999 VTS_03_0.VOB
-r--r--r-- 1 4294967295 4294967295 1073565696 Aug 25 1999 VTS_03_1.VOB
-r--r--r-- 1 4294967295 4294967295 1073565696 Aug 25 1999 VTS_03_2.VOB
-r--r--r-- 1 4294967295 4294967295 1073565696 Aug 25 1999 VTS_03_3.VOB
-r--r--r-- 1 4294967295 4294967295 1073565696 Aug 25 1999 VTS_03_4.VOB
-r--r--r-- 1 4294967295 4294967295 1073565696 Aug 25 1999 VTS_03_5.VOB
-r--r--r-- 1 4294967295 4294967295 1073565696 Aug 25 1999 VTS_03_6.VOB
-r--r--r-- 1 4294967295 4294967295 983531520 Aug 25 1999 VTS_03_7.VOB
d3:/home/glmull#
```

As you can see, the files containing the actual video are going to be the VTS_03_X.VOB files, which are each 1 GB or so in size. And there are 6 of them, X=1,2,3,4,5,6. (7 may also be one of them, I don't remember.) In order to play the movie with MPlayer we would need to play each individual VOB file in sequence. A typical command line would be:

```
mplayer -vo xv -fs -dvdauth /dev/dvd -aid 128
/mnt/hdc/VIDEO_TS/VTS_03_1.VOB
```

I didn't try this but instead of me having to get up out of the bed and type in the command for the next VOB file, why not create a script that has the following commands in it and run the script?

GLENN'S INVENTION

```
#!/usr/bin/perl $begin=1;
$end=6; for($i=$begin;$i<=$end;$i++) {
$play_vob="mplayer -vo xv -fs -dvdauth /dev/dvd
-aid 128
/mnt/hdc/VIDEO_TS/"."VTS_03_".$.i.".VOB";
system($play_vob); }
```

I haven't tested the above script but I do intend to. The idea is to have MPlayer play through each VOB file in sequence, which I would imagine is a trivial thing to have it perform. So if that works, then here's another player that will play through an entire dvd movie although we have to wonder would it cause a pause or other noticeable interruption during the transition from coming to the end of one vob file and beginning to play the next on in sequence.

The following is for the dvd movie disc called "Creature".

```
glnull@d3:~$ mount /mnt/hdc
glnull@d3:~$ mount
/dev/hda1 on / type ext2 (rw,errors=remount-
ro,errors=remount-ro)
proc on /proc type proc (rw)
/dev/hda3 on /usr type reiserfs (rw)
/dev/hdc on /mnt/hdc type udf
(ro,noexec,nosuid,nodev,user=glnull)
```

As you can see above, once we put the disc in the drive and mount it, we can verify that it is mounted as a udf filesystem. Then we can examine the contents of the disc, for whatever reason we might want to do that.

```
glnull@d3:~$ ls -l /mnt/hdc
total 4
dr-xr-xr-x 2 4294967295 4294967295 40 Jul
10 2000 AUDIO_TS
dr-xr-xr-x 2 4294967295 4294967295 560 Jul
10 2000 VIDEO_TS
```

The directory structure is always the same, with the above two directories at the top level.

With the above disc, we can see that VTS_07_X.VOB where X=1,2,3,4 is where the movie is stored. The other VOB files, which are much smaller are maybe trailers or other short "featurettes" that the dvd has but they aren't part of the movie probably. But we can play (or attempt to) them all with MPlayer to find out what they are.

In order to watch the movie, we just need to play the VTS_07_X.VOB files in order, from 1 to 4 I guess. Whether we have to manually play each VOB file or whether the software can figure it out is a question that will have to be answered by each software program individually. Some do and some don't. Notice the filesizes. The VTS_07_X.VOB files are all in the 1 GB range. There are 4 of them. So we can imagine that the dvd disc holds in the range of 4 GB of stuff at least.

The above disc is the movie called "MARS" which was a pretty nice movie, kind of reminiscent of "Total Recall" but a little more realistic. It is about a company that is mining fuel off Mars. I watched and enjoyed it using Videolan's vlc program (videolan client). vlc can play through all the VOB files seamlessly without me having to play each one individually. Further, I don't even have to know which VOB files are the ones that need to be played. vlc figures it all out. Basically, I can put the disc in and type "vlc" and that's all I have to do (other than clicking on play button on the nicely done user interface).

"The Langoliers" is a movie you've probably seen. Here is what the dvd contents look like when I mount that disc:

```
d3:/home/glmull# ls -l /mnt/hdc/VIDEO_TS/
total 7294616
-r--r--r-- 1 4294967295 4294967295 16384 Aug 25 1999 VIDEO_TS.BUP
-r--r--r-- 1 4294967295 4294967295 16384 Aug 25 1999 VIDEO_TS.IFO
-r--r--r-- 1 4294967295 4294967295 159744 Aug 25 1999 VIDEO_TS.VOB
-r--r--r-- 1 4294967295 4294967295 26624 Aug 25 1999 VTS_01_0.BUP
```



```

glmull@d3:~$ /sbin/lsmmod
Module      Size      Used by
udf          80816     1
sb           7456      0
sb_lib      34032     0 [sb]
isa-pnp     28656     0 [sb]
uart401     6416     0 [sb_lib]
opl3        11360     0 (unused)
sound       56992     0 [sb_lib uart401]
opl3]
soundcore   4080      4 [sb_lib sound]
ide-cd      26464     1
cdrom       27392     0 [ide-cd]
glmull@d3:~$

```

Don't forget, if you don't have udf filesystem support in your kernel or loaded as a module, nothing is going to work! But let's push on. Let's take a look at the VOB files on this disc:

```

glmull@d3:~$ cd /mnt/hdc/VIDEO_TS/
glmull@d3:/mnt/hdc/VIDEO_TS$ ls -l
total 4272458
-r--r--r-- 1 4294967295 4294967295 6144 Jul 10 2000 VIDEO_TS.BUP
-r--r--r-- 1 4294967295 4294967295 6144 Jul 10 2000 VIDEO_TS.INFO
-r--r--r-- 1 4294967295 4294967295 49152 Jul 10 2000 VTS_01_0.BUP
-r--r--r-- 1 4294967295 4294967295 49152 Jul 10 2000 VTS_01_0.IFO
-r--r--r-- 1 4294967295 4294967295 1146880 Jul 10 2000 VTS_01_0.VOB
-r--r--r-- 1 4294967295 4294967295 1073709056 Jul 10 2000 VTS_01_1.VOB
-r--r--r-- 1 4294967295 4294967295 1073709056 Jul 10 2000 VTS_01_2.VOB
-r--r--r-- 1 4294967295 4294967295 1073709056 Jul 10 2000 VTS_01_3.VOB
-r--r--r-- 1 4294967295 4294967295 1073709056 Jul 10 2000 VTS_01_4.VOB
-r--r--r-- 1 4294967295 4294967295 78903296 Jul 10 2000 VTS_01_5.VOB
glmull@d3:/mnt/hdc/VIDEO_TS$

```

So we can see that the VOB files that contain the actual movie are VTS_01_X.VOB, where X=1,2,3,4 in that order. If we play them in that order, we will see the entire movie. What about VTS_01_5.VOB? What is it? Well, you could use MPlayer to play it if you wished, in order to see what it was. This is a pretty simple disc, it seems. Now let's watch the movie.

OMS

<http://www.linuxvideo.org/oms>

I had some problems installing oms but eventually I managed to get it installed. omi_gtk had some weird output and stuff when I first tried running it. It asked me a question about dvdbb or something and said do I want to go download it off the net because it couldn't find it in the cache? Well, I hit the return key and what do you know? I get a video window and the DVD begins to play. I am not sure if omi_gtk plays back the video seamlessly although I suspect that it does but I didn't go through an entire movie because it suffered from an audio/video sync problem common to some of the other programs - the half second syndrome - where the audio and video are not in sync but stay the same distance apart throughout the entire movie. There is no drift, just a constant temporal difference of about .25 to .5 seconds but that's enough to really make it hard on the movie watcher. You can't really enjoy the movie when you have that problem.

This would be a nice program to run if only the audio/video sync problem would not exist. The interface is a nice slick looking one although it doesn't have much meat on it. Just a play, pause, stop and similar buttons. But I guess that's all you really need. You don't have to know anything about the names of the VOB files that house the actual movie either. So omi is very userfriendly. A newbie could put a dvd in their drive, run omi_gtk from a command line and then click on the play button and be watching the

movie from the beginning.

That's pretty nice. But on my system, it isn't nice because of the audio/video sync problem. I am not absolutely certain of it but prior to setting the region code using the regionset utility (discussed earlier), I don't think the program would play dvd movies. After I set the region is when I think it would allow me to play them. So it is very important to set the region code on some dvd-rom drives.

<http://www.linuxvideo.org/oms>
<http://linuxvideo.org/docs/wiki/moin.cgi/OmsHowTo>
o

XINE

I've heard some really good things about xine so I was expecting good things coming in. In fact, this is the first dvd playing program that I downloaded, installed and tested. But I was really disappointed and still am. This is the only player out of all the players discussed herein that I didn't get video output from.

That's pretty sad but that's the truth. I tried installing xine_complete and it didn't work at all. I couldn't get anything. There was a problem with CSSAUTH or something. So I tried just installing the normal xine program and patching in the d4d plugin. I think things almost worked there but they still didn't work. I can get some gibberish on the screen and I can see just a little of the video but only for a few seconds and then things bomb out with ac3 errors. The program can play my video cds fine. So I believe I installed it correctly but for some reason it isn't working properly.

The user interface for xine is pretty slick looking but unfortunately, I wasn't able to watch dvds with it, except for one dvd that I have which would play. I suspect that disc is not CSS encrypted but I am not sure how to verify that.

The main page for Xine is at <http://xine.sourceforge.net> but you won't find a CSS-enabled Xine there. What you can do though is download the non-CSS version of Xine and then patch it with the d4d plugin. Patch is really a misnomer, as the d4d patch simply installs into a library directory and doesn't really touch the source codes of the xine sources.

There are basically two ways to get a css-enabled xine. You can download the non-css version of **xine-lib-current.tar.gz** and **xine-ui-current.tar.gz** from <http://xine.sourceforge.net> and compile and install them and then compile and install the **xine_d4d_plugin-0.2.x.tar.gz**. The d4d plugin is available at http://www.geocities.com/captain_css/ In fact, it looks like there has been some recent updating of the plugin so maybe I can retry xine in the near future and see if I can watch some of my css dvds.

The Xine homepage at sourceforge has some links that will allow you to find everything you need, relating to xine. For example, do you want to get a pre-patched version of **xine-lib-current.tar.gz** and **xine-ui-current.tar.gz** so that you don't have to bother with downloading the d4d patch? You can do that [here](http://gape.ist.utl.pt/ment00/linuxdvd.html): **<http://gape.ist.utl.pt/ment00/linuxdvd.html>**. It's called **complete_xine_0.4.3.tar.gz** or whatever the latest version is.

On the dvd that xine did work on for me, I noticed that the audio and video seemed to be in sync, which I appreciated. So hopefully, I'll be able to get xine working on my system because it seems like a really nice program with a nice user interface.

The following is typical of the error message that I get when trying to play some of my dvd discs:

```
input_d4d: unable to open raw dvd drive
(/dev/rDVD);
No such device
load_plugins: input plugin found : d4d
<strong>load_plugins: cannot open input plugin
/usr/local/lib/xine/plugins/xineplug_inp_dvd.so:</strong>
<strong>/usr/local/lib/xine/plugins/xineplug_inp_dvd.so: undefined symbol: CSSAuthDisc</strong>
load_plugins: input plugin found : stdin_fifo
load_plugins: input plugin found : TCP
load_plugins: input plugin found : file
load_plugins: input plugin found : RTP
load_plugins: input plugin found : VCD
load_plugins: demux plugin found : MPEG_TS
load_plugins: demux plugin found : MPEG
```

I have a dvd disc entitled "Thinner" (Steven King) that Xine will play, even the non-css version of xine so that means my problem with xine is with css-enabled discs. Xine does a good job on that disc though. Smooth video, audio/video in sync.

VIDEOLAN

<http://www.videolan.org>

Videolan wins the top spot for ease of use and quality of dvd playback, in my testing. The dvd playback is smooth and the audio is in sync with the video. Another nice feature of the program is that you don't have to mount the dvd disc or specify a path to the particular VOB file you want to play - the program figures it all out and all you have to do is hit the play button. It's really about that simple. The documentation says that you need gnome but I guess you really don't, it can use a gtk+ interface and that is the interface that I used. The interface, by the way, is pretty nice. It's not fancy but it gets the job done and lets you do alot of things. It's really nice to use videolan's player and that is the player that I will use because all of the other programs either have a problem with audio/video synchronization or, as in the case of Mplayer, don't support seamless integration and playback of multiple VOB files. It's nice not to have the movie stop in 3 or 4 places and have to go back to the computer and load the next VOB file and play it. Videolan plays back widescreen format just fine and of course, it also plays back

normal aspect ratio too. So it's a real nice package.

Unfortunately, when running videolan, it seems every 30 minutes or so, the screen goes blank and the monitor seems to lose the video signal. What I do is go and wiggle my mouse on the mousepad and click some buttons and the screen comes back up but that's a real hassle. I am not sure if this is a problem with any of the other dvd playing programs. However, since vlc has such a nice user interface, all I have to do is move the slider bar backwards just a bit and I'm where I left off.

MPLAYER

<http://www.mplayerhq.hu>

MPlayer is a very nice program too, because it happens to support DVD playback now and I gave it a try. What I found was a command line-only tool but that's ok. In fact, MPlayer wins second place in the shootout here. This is a useable program to watch your DVD movies. The video and audio did not have a sync problem at all but if they did I would be able to use the + and - keys on my keyboard to change that synchronization ahead or backward, which is a very useful thing that I wish all the other players would allow because if they did, they would be useful to me.

However, MPlayer is the only program out of all of them that appears to let the user change the temporal relationship between the audio and video just in case your hardware gets them out of sync for some reason. MPlayer gets minuses in the fact that it doesn't support seamless playback of all the VOB files that make up the DVD movie so what happens is, about 3 or 4 times throughout the movie, you'll have to go back to an xterm and retype the command (or use the arrow key to pull back up the previous command and change the name of the VOB file) to play the next VOB file. This requires, of course, that you know which VOB file that is, which of course, requires that you know a little about the mechanics of the DVD filesystem and how to take a look at a disc. Again, that is not too much to ask but certainly an end user might not feel at home doing that and might not want to have to go to all that trouble. But what trouble?

The command line is kind of hard to get ahold of because it has alot of switches and options and the -aid option is very important. Don't ask me what the heck it means but -aid 128 produces the audio correctly. But again, for all its minuses (no gui, unseamless playback, cryptic command line options) you'll find that you get what you pay for except in this case. You get something that you can use to play DVDs.

A typical process to play a dvd movie using MPlayer would be first mounting the disc to see what VOB files were the ones you needed to play (specify on the command line) and then typing the following:

```
d3:/home/glmull# mplayer -vo xv -fs -dvdauth /dev/dvd -aid 128
/mnt/hdc/VIDEO_TS/VTS_03_7.VOB
```

```

The output I get is the following on the console:
MPlayer 0.18pre5 (C) 2000-2001 Arpad Gereoffy <carpi@thot.banki.hu>
Reading /root/.mplayer/codecs.conf: 15 audio & 32 video codecs
font: can't open file: /root/.mplayer/font/font.desc
DVD is encrypted, issuing authentication ...
LBA: 3167318
DVD title key is: DP0C9AC072
DVD auth sequence seems to be OK.
==> Found video stream: 0
Detected MPEG-PS file format!
==> Found audio stream: 128
mpeg2dec-0.2.0-release (C) 2000-2001 Aaron Holtzman & Michel
Lespinasse
libmpeg2: Using MMX for IDCT transform
Libmpeg2: Using 3DNOW for motion compensation
VIDEO: MPEG2 720x480 (aspect 2) 29.97 fps 10000.0 kbps (1250.0
kbyte/s)
[V] filefmt:2 fourcc:0x10000001 size:720x480 fps:29.97
ftime:=0.0334
Detected audio codec: [ac3] drv:3 (AC3)
dec_audio: Allocating 6144 + 65536 = 71680 bytes for output buffer
!! Tibac3 (core v0.6.2) : using 3dNow-dsp! optimization
2.0 Mode 48.0 KHz 192 kbps Complete Main Audio Service
AUDIO: srates=48000 chans=2 bps=2 sfmt=0x10 ratio: 24000->192000
Detected video codec: [mpeg12] drv:1 (MPEG1 or 2)
vo: X11 color mask: PFFF (R:F800 G:7E0 B:1F)
vo: X11 running at 800x600 depth: 16 ("0.0" => local display)
VO: [xv] 720x480 => 800x533 Planar YV12 fs
VO: Description: X11/Xv
VO: Author: Gerd Knorr <krraxel@goldbach.in-berlin.de>
Xvideo image format: 0x32595559 (YUY2) packed
Xvideo image format: 0x32315659 (YV12) planar
Xvideo image format: 0x30323449 (I420) planar
Xvideo image format: 0x36315652 (RV16) packed
Xvideo image format: 0x35315652 (RV15) packed
Xvideo image format: 0x31313259 (Y211) packed
Xvideo image format: 0x0 ( ) packed
using Xvideo port 43 for hw scaling
[xv] dcx: 0 dcy: 0 dx: 0 dy: 0 dw: 800 dh: 600
[xv-fs] dcx: 0 dcy: 33 dx: 0 dy: 33 dw: 800 dh: 533
Disabling DPMS
stat: 1
AO: [oss] 48000Hz Stereo Signed 16-bit (Little-Endian)
AO: Description: OSS/ioctl audio output
AO: Author: A'rpi
ao2: 48000 Hz 2 chans Signed 16-bit (Little-Endian)
audio_setup: sample format: Signed 16-bit (Little-Endian)
(requested: Signed 16-bit (Little-Endian))
audio_setup: sing 46790 Hz samplerate (requested: 48000)
audio_setup: frags: 16/16 (4096 bytes/frag) free: 65536
Start playing...
[xv] dcx: 0 dcy: 0 dx: 0 dy: 0 dw: 800 dh: 600
[xv-fs] dcx: 0 dcy: 33 dx: 0 dy: 33 dw: 800 dh: 533
A:10772.5 (10772.3) V:10772.2 A-V: 0.020 ct: -1.225 41657 19%
0% 4.9% 0
Exiting... (Quit)
Successfully enabled DPMS
d3:/home/glmull# EXIT

```

I guess you can figure that the `-vo xv` option on the command line chooses XVideo output as the rendering method. The `-fs` option makes the video full screen, which I like. The `-aid 128` sets the audio channel to channel 128. I guess there are other channels, for example, a french or spanish. (Type `mplayer --help` for more details.) We assume that `/dev/dvd` is a symlink to wherever your dvd-rom drive really is, such as `/dev/hdc`. And of course, the path to the VOB file is necessary, as MPlayer has no built-in intelligence that would allow it to figure out what VOB files are the ones to play the movie using and in what order.

I am not sure its necessary to have the disc mounted or not. Try it both ways. I think it worked with both ways for me. MPlayer is a nice program, although not as polished and as user-friendly as it could be.

OGLE

<http://www.dtek.chalmers.se/groups/dvd/>

Ogle would be a great package if it only were able to keep the audio and video in sync. I am not sure why it didn't do that on my system but it's what I call the "half-second syndrome" where you see peoples' mouths move but you don't hear the sound until half a second later, or maybe a bit shorter but the end effect is a very troubling scenario where you can't enjoy the movie. I would rather go buy a dvd player than have to struggle through something like that.

Basically, ogle is an interesting player and it does produce good quality video and it let me use the XV extension too for my video output. I tested it out with "The Hard Way" and it performed very well, with the exception of the audio/video sync problem.

Ogle supports menus, where you can access stuff like biographies of the actors, teaser/trailers for the movie, soundtracks and things like that. None of the other products that I tested seemed to be able to access the menu so apparently, this may be a first for linux dvd players, as the ogle homepage suggests.

However, the menu feature doesn't work perfectly, as I found out with my "MARS" DVD. On a DVD player, a standalone device, it would work perfectly. So there is some work left to do I guess but hey, it's a good start and well-appreciated.

OPTIMISING

The Xine FAQ contains some good information about optimizing your system, in general for dvd movie playing. It doesn't matter what dvd software you use, this FAQ can probably help alot. The XVideo extension is something you might want to try and use and the Xine FAQ tells you how to enable that support on your system, in general, and in particular (for some types of video cards, such as the one that I have, a savage4-based chipset).

The Xine FAQ also discusses raw devices and tells us that setting up the dvd drive to be accessed as a raw device will improve performance. Well I tried those instructions but it seemed to slow my frame rate down and I couldn't enjoy the movie. I am not sure why that happened but it seemed to happen. So I didn't use raw device access in my further testing, I just used normal access, which seems to work fine with no performance issues running at 16-bit color. I'll probably try and test raw devices further in the future though.

I guess optimization can make a difference but in the end, you have to do what works for you and your hardware. If there doesn't seem to be a performance problem then you might not need to set up raw devices and use MTRRs and things like that but I guess it wouldn't hurt to learn about them and try and set your dvd drive using dma if it isn't already (use the `hdparm` utility, I guess!).

The Xine FAQ is installed into `/usr/local/share/doc/xine/` on my system when I install xine. So I can read it there.

CONCLUSIONS

Playing DVD movies on your computer should be a right. You paid the money for the product and you should be able to watch it, using hardware that you paid for. Certainly the DVD people that have the most to lose from CSS being broken are the same people

that could realize that if people can play DVDs on their linux box, they will go buy them. If they can't then are they going to go buy a standalone player?

By allowing linux users to play DVD movies on their computer, the industry is really going to benefit financially. The interesting thing is that most of the software mentioned in this article is under the GNU GPL, it's open source. Even stuff that may have questionable legality under existing legal doctrine in the United States. This is an important lesson for people to consider.

Just because something is open source does not mean that it could not have other problems with it, such as patent issues or other types of issues such as circumventing encryption protection mechanisms. That is why for example xine doesn't ship with css stuff inside. You have to patch it or get the modified build from a non-official xine website source because the official site recognized the questionable legality of what the css stuff does, which is that it decrypts dvd movies and allows them to be played.

The conclusion I have come to is that I can use Linux to play my dvd movies. And that means that when there is a good movie out, I might consider purchasing it because I have something that can play it.

RESOURCES

<http://dxr3.sourceforge.net>
http://directory.google.com/Top/Computers/Software/Operating_Systems/Linux/Hardware_Support/DVD/
<http://linuxtv.org/dvd/>
<http://linuxtv.org/dvd/links.xml>
<http://dara.notbsd.org/~aholtzma/ac3/main.php>
<http://www.opendvd.org>
<http://dvd.sourceforge.net>

Slashdot articles:

<http://slashdot.org/article.pl?sid=00/04/03/115218&mode=flat>
<http://slashdot.org/article.pl?sid=00/03/31/156238&mode=flat>
<http://slashdot.org/article.pl?sid=00/04/12/0826247&mode=flat>
<http://slashdot.org/article.pl?sid=99/10/30/2321235&mode=flat>
<http://slashdot.org/article.pl?sid=99/11/02/1251243&mode=flat>
<http://slashdot.org/article.pl?sid=99/10/26/0810224&mode=flat>
<http://slashdot.org/article.pl?sid=99/10/20/0042238&mode=flat>
<http://slashdot.org/article.pl?sid=99/09/02/1426205&mode=flat>
<http://slashdot.org/article.pl?sid=99/07/15/0533246&mode=flat>
<http://slashdot.org/article.pl?sid=01/06/19/1343253&mode=flat>
<http://linux.com/newsitem.phtml?sid=1&aid=11799>
This is a great article entitled "Playing DVDs with

Xine".

<http://www.ox.compsoc.net/~swhite/DVD/>
<http://www.cs.cmu.edu/~dst/DeCSS/Kesden/>
<http://www.ils.unc.edu/gbnewby/DVD/DVD-Playing-HOWTO.txt>

IN RETROSPECT

The "In Retrospect" section will be used for updating things or making further comments about the article. Sometimes, people email me with some useful information. I will put some of that type of information here, so that other people can find it. If you have some useful information that you would like to share, please feel free to write in.

Date: Thu, 4 Oct 2001 02:30:15 -0400
From: Bradley M Alexander <storm@tux.org>
To: glmull@machineofthemoon.org
Subject: DVD article
Glenn,

Thanks for the well-written article on dvd playing that you wrote. I had questions about which of the players was best, which you answered rather conclusively.

I have a couple of questions based on the article. 1, like you, tried most all of the X-based dvd players and I have yet to watch a DVD on my Linux box. Fortunately, I have a dual-boot Mac Powerbook that I can watch DVDs on (its the only reason I keep MacOS around...)

Let me give you my system specs and situation, and hopefully you will have a solution. My machine, defiant is as follows:

```
1GHz Athlon-C (266MHz FSB)
Asus A7V133 Motherboard
640MB PC133 RAM
30GB Maxtor HD (on /dev/hda)
30GB Quantum HD (on /dev/hdb)
Matshita SR-8586 (on /dev/hdc, using IDE-SCSI
emulation, so /dev/sr0)
LG CD-RW CED-8120B (on /dev/hdd, also IDE-SCSI,
/dev/sr1)
SBLive Platinum 5.1
nVidia GeForce 2 Pro/64MB
```

Softwarewise, I'm running Debian/GNU Linux (sid/unstable), and I am running devfs (which is also why I haven't set up /dev/dvd permanently yet). I'm running kernel 2.4.9.

As I said, I have tried vlc, xine, ogle and OMS, and have gotten none of them to work. With vlc, which I just tried again based on your article, when I try to run as my user, I get

```
[storm@defiant lib]$ vlc
VideoLAN Client - version 0.2.83 Ourumov -
(C)1996-2001 VideoLAN
Inconsistency detected by ld.so: dl-deps.c: 495:
_dl_map_object_deps:
Assertion `map->l_searchlist.r_list[0] == map'
failed!
```

However, when I run it as root, it seems to work. I had to upgrade libdvdcss to 0.0.3, but when I fire up vlc, and click disc, if I change the device to either /dev/hdc (which, in devfs, doesn't exist because of ide-scsi), /dev/sr0, or /dev/scsi/host0/bus0/target0/lun0/cd, I can hear

the drive spin up. But when I click play, nothing happens. Unfortunately, I have had no luck in selectively employing ide-scsi.

Can you give me any ideas about what I'm doing wrong?

Regards,

--

--Brad

Bradley M. Alexander, CISSP | Co-Chairman,
Beowulf System Admin/Security Specialist |
NoVALUG/DCLUG Security SIG |
Debian/GNU Linux Developer |
storm@debian.org |
| storm@tux.org

You can't have everything...where would you put it?
If anyone has any suggestions to help Bradley, please email me or him and we would appreciate it.

Regarding this article:

<http://www.machineofthemoth.org/articles/a53/index.html> you mention that your screen would blank out every 30 mins requiring you to move your mouse to get it back. The problem may be your X settings. If you are not familiar with the command 'xset' take a look at that. 'xset q' should show you your settings. The relevant options are the DPMS (turn that off, 'xset -dpms'), screen blanking ('xset s noblank') and standby mode ('xset dpms 0 0 0'). Of course, checking out the man page will also give you lots of info. I had a problem with my mandrake 8 box blanking out on me during DVDs until I messed around with xset. Hope that helps...

Erick Waldchen

Global Computer Supplies

I will try this and see if it helps with vlc.

Looks like MPlayer-0.50pre1 which can be downloaded from the MPlayer website at:

<http://www.mplayerhq.hu/homepage/dload.html>

allows the user to not have to specify the vob file to play at the command line. At the time I wrote this article, I don't think 0.50pre1 was out and so I was using 0.18pre5, which I believe did require the long command line specifying the vob file.

But now, version 0.50pre1 takes care of that issue. I downloaded, compiled and installed and when I type mplayer -dvd 1, that will play my dvd, assuming that /dev/dvd is a symbolic link to the dvd-rom drive, be it /dev/hdc or /dev/hdb. In addition, the good thing about mplayer on my system is that it doesn't seem to have the screen go blank and I have to go and type on my keyboard or wiggle my mouse to get the screen back (videolan had this problem in my testing). So I guess if I had to update my ratings, I might just put MPlayer at the top.

As several kind readers have pointed out, and I have verified in my testing, MPlayer can play an entire dvd movie without you having to know anything about

dvd VOB files or anything. Just pop the disc in and type mplayer -dvd 1 or mplayer -dvd 2 or whatever title you are trying to play on the dvd.

Of course, MPlayer still does not seem to come with a gui, although they say one is in beta but I don't really feel like I'm missing out because I can do alot from mplayer's command line. Type mplayer -h for all the options that are available.

A reader writes in with some news on mplayer:

Hi,

I've just read your article about linux dvd players, including mplayer. You used teh very old 0.18pre5 version, so I decided to tell you 0.50pre is out. It has "real" dvd support, using IFO parser instead of playing individual VOB files... Oh, and now it has a GUI too (optional), but it's in early beta stage yet.

A'rpi / Astral & ESP-team

--

mailto:arpi@thot.banki.hu

<http://esp-team.scene.hu>

Stephen Moore writes in that he enjoyed my article:

It was fun for me also getting dvds to work on linux, I started playing dvds on linux and then moved to buying a dvd player (and 5.1 receiver and speakers... so my \$30 dvd on linux cost me a lot of money)

I use xine, I found oms unusable I tried vlc but just preferred xine. I have never experienced the 1/2 second lag with xine ever I currently am running 0.9.1 and always compile it myself from source, not because I dont like rpms, but because I want it now!

Anyhow its probably worth perservering with xine as it now does ac3 passthrough as well as deinterlacing. Cheers

--

Greg Breland writes:

Very good article, nice overview of all the DVD players for Linux. I started with OMS and could not get it to install and then tried Xine. It worked so well, I have never tried any of the others.

Your audio problem with Xine(and probably the other players) is cause because of the fact that the OSS sound driver don't support realtime playback for most cards. I had the same problem you are experiencing and I fixed it by using artsd instead of esd or oss. esd is apparently horrible at realtime sound. With OSS, my audio would get noticeably off every 30 minutes and I would have to stop and restart play. With ESD, audio was off from the very begining. Now when I want to watch a DVD, I just kill esd and start artsd and there are no problems. You can also use alasa, but I have not tried that yet even though I have alasa installed.

I am glad you mentioned WM and DVD playback problems. I was running KDE 2.2.x and my

framerates were horrible. kdeinit was using 30% of my processor for some reason while playing DVDs. I swithed to gnome and sawfish and have had no troubles since, even though I still use artsd for audio.

BTW, my machine is a 550mhz dual celeron with 256mb RAM and a Matrox G400 dual head with a 21" and 19" monitors. I can play a dvd, surf the web, and compile PHP/Apache without dropping a frame using Xine 0.9.1, so I think your machine should be able to caculate weather forcasts while playing a DVD with no problem. 0.9.0 did have some problems though, so maybe that is related.

You should also mention that you MUST have DMA enabled on your DVD and Xv extensions loaded to get smooth playback. Make sure you read the stuff xine prints to the command line as it will tell you a lot about what is not optomial with your system, like the lack of realtime audio.

Lastly, DVDNAV is pretty cool if you like menus. It is still alpha, but worth messing with for the cool factor.

Greg Breland
<http://www.mozillanews.org>

To turn on DMA, I use `hdparm -d1 /dev/hdc` and that does seem to make a difference with Xine. Replace `/dev/hdc` with whatever device your dvdrom drive is on.

There are a few things that I really should point out. The window manager, apparently, can affect the quality of your dvd playback. How do I know? I first tested some of the programs with the blackbox window manager. But my frame rates were low and the motion of the video was nowhere near as fluid as it should have been. I could not enjoy the movie that way. Eventually, several days later, I guess I was in the mood to try a different window manager and so I put `fvwm95` on my system. And guess what? The dvd playback was much better, actually useable. So guess what? I stuck with `fvwm95`. If at first you don't succeed...try another window manager. Another thing that has recently popped up is the fact that Videolan's vlc client may not work for all dvd movies as well as it does for some of the ones that I tested out.

Unfortunately, the Steven King "Thinner" DVD, is an example of where I needed a command line tool that would allow me to specify the exact VOB files to play, and in what sequence. Hopefully this is something that doesn't happen often but putting the disc in and typing "vlc" doesn't do the right thing. It plays a trailer and then stops. Hmm. But MPlayer can handle that disc with no problem. Another thing I need to say is that I tried out a later version of xine and I managed to get it working with css-encrypted discs. I downloaded the non-css version of xine from the official xine website and patched it with the d4d plugin. This works. The problem is it suffers from the half-second syndrome common to many of the other programs that we looked at in the article. That's really

a disappointment because otherwise, xine is a very nice program and probably would be the top dog. The video is smooth and the interface is about the best I've seen on a dvd player for linux. Anyway, just my 2 cents, in retrospect. Also I guess I should admit that OMS seemed to have low frame rates on my system, even though I didn't mention it in the main article. Yes, it's true, it's a little on the unusable side for me because of non-fluid video. I think that's a rather common complaint of it though. For the other problems, the audio/video sync problem I would be interested in knowing it that were a hardware problem on my system or something that other people have encountered as well.

Another thing worth mentioning is I have onboard sound. I also have a soundcard, an ISA ESS1868 one. I tested dvd playing out with both. The onboard sound apparently caused low frame rates for the video, such that watching and enjoying the movie became impossible. I don't know why this is except, as one person suggested, decoding the audio is occuring using software rather than hardware but on a 1.2 Ghz system with 1 GB of RAM and a 10X dvd drive? But on my trusty ISA soundcard, no problems. It's a difference between night and day, as far as the video frame rate and quality of fluid video go. Unuseable versus useable.

So don't give up. You'll get there. One thing I didn't cover in the article is the process of finding a dvd-rom drive and installing it on your system and setting it up with linux. The reason why I didn't do that is it's not very difficult and it's pretty much just like putting a cd-rom drive in. All you have to do is register it in `/etc/fstab` and you're good to go. A dvd-rom drive can be used as a cd-rom drive too so it will read all your cdroms too. And we didn't discuss prices of dvd-rom drives but let me tell ya. You don't need to go without. You can get one for \$30 maybe. A brand new one. Isn't that a great deal? Places like your local neighborhood Best Buy sometimes have deals where you can get them for \$40, which is not bad either. My experience with cd-rom drives has been that they all work, if they are IDE. So I would imagine that most dvd-rom drives that are IDE will work too.

As I mentioned, I managed to get Xine working. But as I discussed earlier, there are two ways to get a css-enabled Xine. The first way, downloading the official non-css lib and ui tar.gz files and compiling and installing and then compiling and installing the d4d patch, that worked for me using version 0.9.0. However, the 2nd way still doesn't produce the results I am looking for. The lib and ui files that are already patched, those still don't let me play my discs, I seem to get errors. Not sure why but it happens. This just goes to show you that the more ways you have of doing something the better your chances of getting one of them to work. And that's what it's all about, getting it to work.

This article is re-printed with permission. The originals can be found at:

<http://www.machineofthemoth.org/articles/a53/index.html>

Using XML and XSLT

Author: Egon Willighagen <egonw@sci.kun.nl>

Abstract

This article contains the presentation given at the Libre Software Meeting in Bordeaux in July. It explains the XML database used for automatic generation of the LinuxFocus.org(/Nederlands) web site.

INTRODUCTION

The system used for document and translation management in the LinuxFocus project consists of several ASCII files, including *resdb.txt*, *issuedb.txt* and *maindb.txt*.

These files have a fixed format, and they're used to generate web pages. However, they are difficult to extend, and the separated nature of the data makes it hard to manage all the information available for an article.

LinuxFocus did not automatically generate much web content when I started the new database. As an editor on the Dutch team, I was eager to have the index.html files on the web site dynamically generated.

Editing several HTML files each time a new article was translated took a lot of effort and caused many broken links. Therefore, I wanted a new system to which I could add information easily, and from which I could easily generate index pages for the web site. I started working on it sometime in the summer of 2000.

The choice for XML was a bit arbitrary. Suggestions had been made to use a relational database, but I was experienced in XML and preferred a system of text based files. It soon turned out that a new numbering scheme would be useful, because the database could then use one type of ID instead of the two or three schemes then in use. Guido Socher did all the renumbering, which was quite an effort (my thanks!).

The Document Type Definition (DTD) was already in development, and a little bit of content was in the database for testing purposes. With the new uniform numbering scheme, the time was right to load the database with content. After having added about 20 articles, it became clear that this was an enormous project. Writing scripts to use the old files was possible, but not all information that the new database could contain was available, and, as explained, the information that was available was distributed over several files. Fortunately, Floris

Lambrechts got involved, and I have to thank him deeply for adding most of the content to the database. Without his help, the system would not be what it is today.

Along with the new format also came the ability to add new information. And over the past year several new kinds of data have been added to the database. Early extensions were a table of authors, translators, editors and other people involved in LinuxFocus, and file locations. The reason for addition of the latter was that there were several filenaming schemes used since the beginning of LinuxFocus. During the renumbering it was reduced to two schemes. Some files used server side includes and used the .shtml extension, where older articles used .html extensions. The <file> tag can be used to overwrite a default. (The current default uses the format "article" + article number + ".shtml". This might include an optional ".meta" in case the file is in LinuxFocus' meta format.)

Now that the database had reached critical mass, I finally got around to benchmarking the software I was writing. The current XSLT stylesheets are not the first implementation. It was preceded by Perl based code. But with the growing size of the database, performance became important. The first try was simply not good enough. But before I start explaining the tools, I'll explain the database format.

THE DOCUMENT TYPE DEFINITION

XML, first of all, is a syntax specification for markup languages. XML defines how markup should look. The syntax describes the sequence of characters allowed in "well formed" XML document. It declares that a document has one root element and that an element consists of a start tag, content (text, child elements, or both), and an end tag. These tags consist of a "<" character followed by a name and at the end a ">" character. An end tag must have a "/" just in front of the name. Empty tags, like HTML's
, take a "/" after the name. A start tag may contain attributes, and these also have a specific syntax. XML tags look like these:

```
<greeting>Hello, world!</greeting>
```

or for an empty tag

```
<br/>
```

Besides syntax, languages also contain semantics. This describes how certain elements relate to each other. The semantics of HTML declares that the <body> tag should be contained by the <html> element, and not the other way around. The semantics also describe that the element is empty, as is the
 element. If these semantics are given in a formal notation, they can be parsed with a program and used to validate the document using those semantics. One of these formal notations is called Document Type Definition, or DTD for short. If a document passes the validation process, it is called a valid document. You have to be careful with XML

because its validation is very strict.

Now that we know what a DTD is, let's have a look at the LinuxFocus XML Database DTD. For several of the specifications we will provide an example. By examining these examples you will get an idea on how the information is contained in LinuxFocus' XML database.

<DATABASE>

The root element in the LinuxFocus XML database, or one of its extensions/localizations, is the <database> element.

```
<!ELEMENT database (themes?, persons?, issues?,
articles?)>
```

First, note that the "?" means the child element may occur zero or one times. Thus, the database may contain information about LinuxFocus' themes, persons, issues and articles. Since this is very straightforward, I'll move on to a more interesting example.

<THEMES>

The themes are contained within the <themes> element which is a child element of <database>. Each theme has a unique ID, a title, and optionally an abstract and an image.

```
<!ELEMENT themes (theme+)>
<!ELEMENT theme (title*, desc?, img?)>
<!ELEMENT title (#PCDATA)>
<!ELEMENT desc (#PCDATA)>
<!ELEMENT img (EMPTY)>
```

Some of these elements must have attributes. These are also given in the DTD. Any textual content is contained in an element with the xml:lang attribute. The value of that attribute may be any token conform the ISO 3166 standard for country codes. Examples are "en", "fr" and "nl". Both the id and xml:lang attributes are specified in the original XML specification and are part of the XML syntax.

```
<!ATTLIST theme id ID #REQUIRED>
<!ATTLIST title xml:lang xml:lang #REQUIRED>
<!ATTLIST desc xml:lang xml:lang #REQUIRED>
<!ATTLIST img src CDATA #REQUIRED>
```

An example database might look like this:

```
<database>
<themes>
<theme id="hw">
<title xml:lang="en">Hardware</title>

</theme>
</themes>
</database>
```

<ISSUES>

Issues are contained in the <issues> element. Like themes each issue has a unique ID.

```
<!ELEMENT issues (issue+)>
<!ELEMENT issue (title+, published?,
file*)>
<!ELEMENT title (#PCDATA)>
<!ELEMENT published (EMPTY)>
<!ELEMENT file (#PCDATA)>
```

The element <published> flags published issues. The next issue and the SomeLanguage2Eng pseudo issues do not have this element. The <title> element has again the @xml:lang attribute. The <file> element denotes the directory in which this issue is located. It must not point to the index.html, because it is used to determine file locations.

An example (note that we use the @code attribute for sorting):

```
<issue id="ToBeWritten" code="999996">
<title xml:lang="en">Not yet written articles</title>
</issue>
<issue id="September2001" code="200109">
<title xml:lang="en">September2001</title>
</issue>
```

<PERSONS>

Information about authors and translators are stored in <person> elements. Each person must have a unique ID.

```
<!ELEMENT persons (person+)>
<!ELEMENT person
((name|email)*, (homepage|nickname|desc|team)*)>
<!ELEMENT email (#PCDATA)>
<!ELEMENT name (#PCDATA)>
<!ELEMENT homepage (#PCDATA)>
<!ELEMENT nickname (#PCDATA)>
<!ELEMENT desc (#PCDATA|&#x2013;els;)*>
<!ELEMENT team EMPTY>
```

Each person can have the following information: a name, an email address (or more), homepage(s) and nicknames. If the person is also part of a translation team, we add a <team> element. For example, the following line in the <person> element means that Floris belongs to the Dutch team <team xml:lang="nl"/>. Finally, each person can have a description, which may contain additional web links.

An example:

```
<person id="nl-ew">
<name>Egon Willighagen</name>
<email>egonw@linuxfocus.org</email>
<team xml:lang="nl"/>
</person>
```

<ARTICLES>

The articles are of course the most interesting part of the database.

```
<!ELEMENT articles (article+)>
<!ELEMENT article (title+,
(file|personref|abstract|issueref|themeref|
nometa|nohtml|translation|proofread)*)>
<!ELEMENT abstract (#PCDATA)>
<!ELEMENT nohtml EMPTY>
<!ELEMENT nometa EMPTY>
<!ELEMENT translation
```



```

      (personref*,
(reserved|finished|proofread)*>
  <!ELEMENT reserved (#PCDATA)>
  <!ELEMENT finished (#PCDATA)>
  <!ELEMENT proofread (personref*,
(reserved|finished)*>
<|ATTLIST article      id          ID
#REQUIRED
#IMPLIED
      xml:lang          NMTOKEN
(article|coverpage)
      type
"article"
      next              IDREF
#IMPLIED
      prev              IDREF
#IMPLIED>
<|ATTLIST file      xml:lang          NMTOKEN
#REQUIRED
      type              (target|meta)
"target">
<|ATTLIST translation from          NMTOKEN
#REQUIRED
      to                NMTOKEN
#REQUIRED>

```

Each article has at least one title; one for each language. The <file> element can be used to give the article's file location, for both the META format and the HTML version (see example below). In cases where no META or HTML version is available, the optional <nohtml/> and <nometa/> elements may be used. Each article can have an abstract. Having the abstract in the database means it can be used to create index web pages.

The <article> element has five attributes: the required @ID, an optional attribute xml:lang to denote the language in which it was originally written, a @type attribute used for cover pages, which are for translation purposes also treated as articles. Finally, two other optional attributes, @next and @prev, which are used to tie articles from a series together.

An article is associated to an issue and to a theme with the <issueref> and <themeref> elements, both having a @href attribute. The value for this attribute must be a unique ID, the ID of the associated issue or theme.

An example:

```

<article id="article206" xml:lang="en">
  <title xml:lang="en">Using XML and XSLT to build
    LinuxFocus.org(/Nederlands)</title>
  <personref href="nl-ew"/>
  <issueref href="nl-ew"/>
  <themeref href="appl"/>
  <abstract xml:lang="en">
This article shows you how parts of the Dutch web
site of LinuxFocus is
generated with XSLT tools from the XML database.
It compares this with
the (very) much slower DOM tools in Perl.
  </abstract>
</article>

```

A localized <article> element looks like:

```

<article id="52">
  <title xml:lang="nl">Enlightenment</title>
  <file
xml:lang="nl">Nederlands/July1998/article52.html</
file>
  <translation from="en" to="nl">
  <personref href="nl-tu"/>

```

```

      <reserved>2000-09-06</reserved>
      <finished>2000-10-04</finished>
      <proofread>
        <personref href="nl-fl"/>
        <reserved>2000-10-04</reserved>
        <finished>2000-10-04</finished>
      </proofread>
    </translation>
  <abstract xml:lang="nl">

```

Enlightenment is een Linux window-manager met uitgebreide mogelijkheden. Dit artikel bespreekt ze, samen met de installatie en de instelling van E. Dit alles is niet voor beginners daar E op het moment nog in beta-stadium verkeert.

```

  </abstract>
</article>

```

Note that this translation is reserved for translation at a certain date, it is done, but also proof-read. In all cases the person who did the work is linked to with <personref> elements.

For all elements, the best tutorial is the current databases itself:

English Database
<http://www.linuxfocus.org/common/src/article206/xml/db/lfdb.en.xml>

Dutch localisation
<http://www.linuxfocus.org/common/src/article206/xml/db/lfdb.nl.xml>

AUTOMATICALLY MAKE WEB PAGES

One of the reasons for creating this new format was to automatically create web indices from it. Now that we understand (?) the database format let's see how we can use it to generate web pages.

First, a bit of history. The first implementation used Perl modules to interface with the database. Though the interface was very clean, the implementation was very slow. The information was contained in an XML container called Document Object Model (DOM). Most implementations for DOM, however, are very slow, at least much slower than the alternative Simple Application interface for XML (SAX).

But if the task is just web page generation a third alternative seems best: XSLT. This is a XML based transformation language. Many XSLT processor currently exist and most programming languages are supported. Some time ago there was a LinuxFocus article on XML::XSLT, one of Perl XSLT implementations. Since the publication of that article, more implementations have emerged, and there are a few that I recommend:

Sablotron
<http://www.gingerall.com/>

Xalan
<http://jakarta.apache.org/>

Xsltproc

<http://xmlsoft.org/XSLT/>

The examples in the remainder of this article will use Sablotron.

An XSLT processor takes two files for input. One is the XML source to transform. The other is the XSLT stylesheet that defines the transformation. For generation of LinuxFocus web pages the following XSLT stylesheets are available:

(The following are all available from:

<http://www.linuxfocus.org/common/src/article206/xml/stylesheets/>)

issues.xslt

This stylesheet generates a list of issues, with their respective articles.

issuetoc.xsl

This one generates the table of contents for a certain issue.

issuetoc_full.xslt

Like the previous, but with more information.

mainindex.xslt

Generates a list of articles with information on the translation status.

previssues.xslt

A list of all issues that have been published.

recently_translated.xslt

The ten most recently translated articles.

rss.xslt

Generates a RSS file with the ten most recently translated articles.

theme.xslt

This stylesheet generates the index page for a certain theme.

themes_index.xslt

Generates an index of all themes.

vertaald.xslt

Shows all translated articles for a certain language.

Note that these stylesheets are not the latest versions. Contact me or one of the editors of the Dutch translation teams to get up to date versions.

To generate the mainindex.html, for example, the Dutch teams runs:

```
sabcmd stylesheets/mainindex.xslt db/lfdb.nl.xml >
../mainindex.html
```

The stylesheets know where the English root database is, and just needs the localized database as XML input. Some sheets need an additional parameter:

```
sabcmd stylesheets/theme.xslt db/lfdb.nl.xml
'$theme=appl' > ../Themes/appl.html
```

The Dutch index.html is also generated from the database, but uses a bit more complex setup. The index.html is made with Guido Socher's lfpagewriter from a set of preprocessed input files. And these preprocessed input files are generated from a set of .pre files such :

```
<H2>Vorige nummers</H2> <p>Dit zijn de uitgaven
van LinuxFocus in het Nederlands:
<ul>
<!-- macro xslt previssues -->
</ul>
<H2>Recent vertaalde artikelen</H2>
< macro xslt recently_translated -->
```

These files are simply HTML fragments with a macro that applies a stylesheet to you localized database. The processing is done with a program called apply_stylesheets.pl which looks for <!-- macro xslt [stylesheet] --> commands and parses the database with that command. Note that the .xslt extension is omitted. Our Makefile contains:

```
%.shtml: %.pre
    @echo "Making $*..."
    @../xml/bin/apply_stylesheets.pl $*.pre
```

The resulting *.shtml files are used by the lfpagewriter script. The stylesheets that are used to generate the index.html are:

```
issuetoc.xslt, previssues.xslt and
recently_translated.xslt.
```

LOCALIZING

To use this system for other languages, you need to do the following:

1. localize the XML database (like lfdb.nl.xml)
2. localize the stylesheets

The second step is a bit unfortunate. In principle only the text in the output needs to be localized, but the stylesheets do not have localization properties yet. This is possible, however, and I would like to see it implemented.

I recommend using a DTD aware XML editor. In Emacs you can, for example, use the psgml major mode. This will give you the ability to validate the document (with nsgmls). This helps a lot in avoiding mistakes. In Emacs you can then also right-mouse-click to see the elements and attributes you can insert on that specific place in the XML file. (Thanks to Jaime Villate for his excellent talk at the LSM conference in Bordeaux this year.)

Another great help is the Dutch localization of the XML database. If you run into trouble you can consult that file. Though the content is mostly Dutch, you can see how the database elements are organized. If that does not help, you can always email me.

Localizing the stylesheets is probably a bit tricky. Text is intermingled with XSLT commands. The latter you must not touch (unless you know what you're doing), in order to preserve its functionality. I plan to have the stylesheets localized in the future which would mean that you only need to edit a file that contains your translations and no XSLT commands, but this is not yet done.

FUTURE PLANS

OK, this should help you to get started. Most things you can copy/paste from the Dutch files. All files are FDL and GPL. In the next year these are my plans with this system:

- localize stylesheets
- add new stylesheets (for `top_authors.html`, `top_translators.html` and others things we like to see as web pages)
- possibly an interface to a daemon based relational database, like MySQL.
- integrate the system with other LinuxFocus tools (like `gettick` etc.)

This article is re-printed with permission. The originals can be found at:

<http://www.linuxfocus.org/English/September2001/article206.shtml>

Logical Volume Managers

Author: Ramon Casha

ABSTRACT

Ever ran out of disk space on one hard disk while you still have space on another? Or wanted to take some free space from one partition and give it to another? Do you wish you could treat two hard disks as if they were one?

Logical Volume Managers can take away many of the headaches traditionally associated with managing your hard disks.

THE PROBLEM WITH THE TRADITIONAL DISK/PARTITION SYSTEM

One of the big headaches of all the operating systems that I know about - including Linux and Windows - is the fact that the end-user has to be aware of having different hard disks or partitions inside the computer.

When computers were operated only by technical

people this was OK - they were perfectly happy to tinker around with partitioning programs, resizing partitions and so on. That time is over.

Most users nowadays don't want to know how many sectors the secondary slave IDE hard disk has, and the thing is, they shouldn't need to. The most they need to know is that somewhere deep within the belly of their computer is enough storage space to store a grand total of "X" megabytes of data. If that fills up they need to add more. That's it.

The situation in most "default" operating systems at present is that hard disk storage shows up as different "hard drives" - which may either be different physical devices or different partitions on one device.

It is really frustrating when, after using your computer for a while, you find out that the `/home` partition is full while the `/usr` partition is half empty, or vice versa (Windows users, think "drive C: and drive D:").

Alternatively, you might need to save a 650Mb ISO image file, but you only have 400MB free on each of two different partitions. Usually you would end up trying to find directories which you can move from one location to the other, or files you can delete, and so on - a tiresome job.

THE ALTERNATIVE

The ideal situation would be if hard disk storage were added and removed more or less like RAM: you simply add a new disk to your "hard disk pool" and get "x" megabytes of additional space on your filesystem. The space gets added to your existing directory tree instead of requiring you to create a new mount point and mount the new hard disk there.

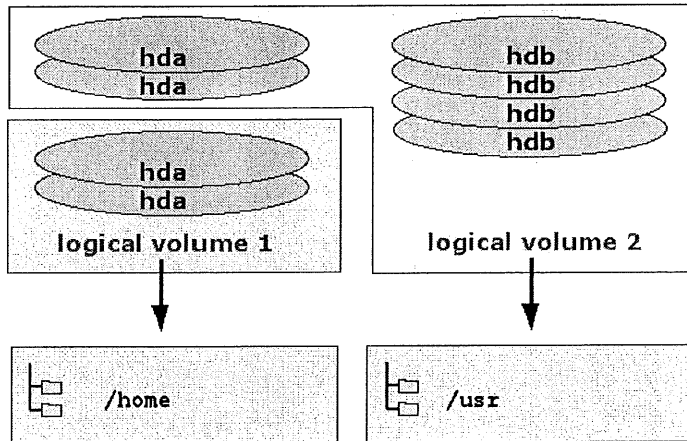
In Linux, this is the job of the Logical Volume Manager. With the LVM, you no longer have to worry about how much space each partition will contain. In fact, you may no longer have to worry about partitions at all.

Instead of partitions and hard disks, you have logical volumes. Here are some of the possibilities offered by the LVM:

- You don't need to accurately calculate beforehand how much space each partition/LV will need. Space allocation can be altered afterwards.
- Several physical disks or partitions can be treated as a single large drive.
- It is possible to add or remove disks from the logical volume. You can remove a failing disk and replace it with a new one - or replace a slow disk with a fast one. Data will be moved off the disk that is going to be removed onto the ones that will remain.
- You can resize logical volumes easily, and can transfer free space from one LV to another.
- You can create a "snapshot" of a logical volume for

hot backups. This allows you to take consistent backups without stopping anyone. The backup will contain the state of the disks at the time the snapshot was started - no matter how long the backup takes.

LVM OVERVIEW



In the above diagram, two hard disks - hda and hdb - are configured as logical volumes. The first hard disk (hda) is divided between the two logical volumes, while the second hard disk (hdb) is completely dedicated to the second logical volume. Thus, assuming the two hard disks are of 2Gb each, you might get a logical volume of 1Gb and another of 3Gb. These are then formatted and mounted as /home and /usr respectively. If, in time, the /home filesystem fills up and the /usr filesystem is not being used completely, you can reallocate some free space from the second logical volume to the first, and it will automatically appear in the /home filesystem. Alternatively, the user could install a new hard disk and add it to the first logical volume, which would effectively add it to the /home filesystem.

Note that the logical volume only replaces the idea of a partition. You must still format the LV to give it a filesystem. You can format your logical volume in the same way that you do a normal partition, using filesystems such as ext2, ext3, ReiserFS, JFS or XFS - all the features offered by these filesystems will be available as usual.

Although there is some impact on performance this is not large. Check out SuSE's whitepaper in the links below - it contains a number of benchmark results.

AVAILABILITY

The LVM module is rather new, and not in wide use so far, so you will probably need to patch and recompile the kernel. Kernels from 2.2.17 can use it with patches, and 2.4 kernels include it as standard (though it might not be compiled in). In any case you will probably want the latest version from the website to work with since it's still in rapid development.

You can expect it to gain acceptance quite fast as stability and support increases, especially if the major distributors such as RedHat, SuSE and Mandrake start offering it as the default configuration for "beginners" installations. For that to happen its installation, configuration and maintenance needs to become more user-friendly, but there are already some distributions which go a long way in that direction.

LINKS

LVM Homepage
<http://tech.sistina.com/lvm/>

SuSE whitepaper - the Logical Volume Manager.
http://www.suse.com/en/support/oracle/docs/lvm_whitepaper.pdf

This article is re-printed with permission. The originals can be found at:

<http://linux.org.mt/article/lvm>

Psionic Portsentry 1.1, the defender of the ports

Author: Georges Tarbouriech <gete@wanadoo.fr>

ABSTRACT

Psionic portsentry is part of the Abacus Project suite of tools (beside portsentry, the suite offers logcheck and hostsentry). It's an IDS (Intrusion Detection System) dedicated to portscan detection and active defense. It works on many Unix flavors including Mac OS X. The main feature of an IDS is to inform the sysadmin about intrusion attempts. Portsentry goes further since it can react on an attack. The latest version (1.1) of this great tool is available from <http://www.psionic.com>, and by the way, version 1.0 is now part of some main Linux distributions (Debian, RedHat...).

WHY USE PORTSENTRY ?

During the last decade, networking grew up at an incredible speed. The goal was to allow communication between machines using different OSes. Thus, proprietary networking systems were slowly replaced with TCP/IP. Internet did the rest ! Today, most networking relies on TCP/IP and TCP/IP relies on ports.

To make it short, ports are attached to programs (clients or servers). The server listens till a client contacts it to establish a connection. The services (the programs above mentioned) are mapped to a specific port. On Unix systems, this mapping scheme is found in the `/etc/services` file. That is, everyone knows which port is dedicated to which service. If everyone knows, crackers know even better ! If we consider a port as a door, when a port is open (listening), it's like an unlocked door. And how do you get into a house ? Usually through the door (unless you prefer the window, it's up to you !). So do the black hats to get into your computer...

The first thing you can do to reduce the risk is to close as many ports as possible, that is to stop services. The less the better. But you hardly can close every port on a networked machine : it wouldn't communicate anymore, that would be a pity !

How to limit the number of active services or how to close the ports is beyond the scope of this article. You'll find a lot of literature on the subject, for instance going to the Linux Documentation Project or searching through LinuxFocus issues (for example, Bastille Linux or Security tools). So for the different ways to protect a machine, a network. On Linux, consider Bastille Linux as a must have.

And this is where portsentry comes. Portsentry can monitor ports and is able to block them if you ask it to do so. It provides you with different operating

modes, some of them being specific to some OSes. As a matter of fact, OSes means Linux.

Portsentry is able to benefit from packet filtering provided with ipfwadm, ipchains or iptables according to the Linux kernel you have. This is true for other Unix flavors using different tools (back on this later). There we are : portsentry greatest feature probably is "auto-blocking".

How does it manage that ? Here we go !

HOW TO INSTALL PORTSENTRY ?

Portsentry comes as a very small tarball. After unpacking the archive, you get source code and various configuration and README files. Obviously, all that stuff is recommended reading. Since we're talking about security, don't forget to download the PGP files (signature and key).

Installing portsentry is obvious : make (your system) and make install. But before installation, you should prepare it. That is, you should read the documentation, change the `portsentry.conf` file to suit your needs and check the Makefile as well as the `portsentry_config.h` file.

For portsentry to be the most efficient, you need at least TCPWrappers. At least, because it would be much better to use it in conjunction with a packet filtering tool. Of course, if you use Linux, you already have everything you need. What about other Unices ?

Since we tested portsentry on different platforms, let's talk about the requirements.

On MAC OS X, it's like for Linux : everything is there. You have TCPWrappers, and ipfw (the BSD version). So just typing "make osx" will do the trick.

On Solaris 2.6 (SPARC), you don't have TCPWrappers. You can get it from:

<ftp://ftp.porcupine.org/pub/security>.

You can also try IPFilter available from:

<ftp://coombs.anu.edu.au/pub/net/ip-filter/ip-fil3.4.9.tar.gz>

(but only if you have the Solaris compiler !!!).

On Irix 6.5, you don't have TCPWrappers either. You can get it as a package from:

<http://freeware.sgi.com/index-by-alpha.html>

For packet filtering you have ipfilterd, coming with Irix but not installed by default.

Last, but not least, unless you find a working `vsnprintf()`, you won't be able to install portsentry on NeXTSTEP... and I didn't find one !

However, once again, you should configure portsentry before installing it. Well, it's up to you, but it should

be better to understand its philosophy before everything else... but it's just a matter of opinion !

HOW DOES PORTSENTRY WORK ?

Fine, thank you !

More seriously, portsenentry relies on configuration files. The most important is the portsenentry.conf file. This is where you tell portsenentry how to react against adversity.

Before dealing with that file, you need to know what the different modes of operations are and what they do. Portsenentry can use six different modes, according to the option you give at startup.

- The first one is "-tcp" and it is the basic mode. With this option, portsenentry will bind the TCP ports found in the config file in the "port configuration" section. It can bind till the limit of 64 ports.
- The second one is "-udp" and does the same as the previous one for UDP ports.
- The third is "-stcp", the "s" meaning stealth. This option and the next ones are only available for Linux. With the "-stcp" option, portsenentry uses a socket to monitor incoming packets, that is the ports are not binding to anything.
- The fourth is "-sudp" and does the same as the previous one for UDP ports.
- The fifth and sixth are "-atcp" and "-audp". These are the most effective options ("a" stands for advanced). Using these options, portsenentry makes a list of the listening ports, TCP and UDP, if you select both, and blocks the host connecting to these ports, unless this host is present in the **portsenentry.ignore** file.

This is much better explained in the README files coming with portsenentry. Accordingly, we won't reinvent the wheel, trying to rewrite the documentation. These README files are compulsory reading (as usual, let's be a bit fascist !).

HOW DOES PORTSENTRY REACT ?

Obviously, portsenentry can generate logs. If you are the kind of sysadmin not reading the logs (shame on you !), you can use logcheck in conjunction with portsenentry. That way, portsenentry sends a mail to inform you about an intrusion attempt. It can write the target host into the /etc/hosts.deny file, for use with TCPWrappers. The local host is able to route the network traffic to a dead host. Last, the local host can drop the packets through the packet filtering local tool. Now you know a bit better about portsenentry philosophy, you can start writing your portsenentry.conf file.

- The first section of the config file concerns the ports. There, you can choose the ports to bind. Remember, this is ignored when using the

advanced modes for Linux. Careful with the ports you want to bind (don't bind port 6000 if you use X, for instance).

- The next section is for advanced detection options. The default is to monitor every port below 1024, either for TCP or UDP. There, you can also exclude the ports you wish to ignore. This is quite important to avoid false alarms, and so huge logging, especially if you have Windos machines in your network. For instance, mind the port 111 if portmap is used and the hosts are not present in the portsenentry.ignore file. Sure, it would be crazy, since portmap means RPC servers like NFS, but you've been warned. That is, NFS is not that secure on Unix machines so don't we talk about Windos machines.

Here comes the usual off-topic section : when I write "Windos", it is not a typo : it is just to insist on the fact Windows is a desktop environment on top of DOS (D for dirty ?)... and it makes one less character to type (how lazy I am !). Sorry for the digression.

- The configuration files section concerns the files used by portsenentry for history purpose or to define the hosts to ignore. If you perform a default install, don't change this.
- The miscellaneous config section allows to turn on or off DNS lookups.
- The response options section is the heart of the configuration. This is where you define how portsenentry will react.

First, are the ignore options. You can choose to block the scans, not to block them or run an external command. Next, in the dropping routes section, you tell portsenentry how to drop the route or how to benefit from the packet filtering tool used on your machine. There are a lot of examples for many platforms. Select the one (and only one !) suiting your needs.

- The TCPWrappers section indicates if you want to write into the /etc/hosts.deny file.
- The external commands section allows to define a command to run when an host connects.
- The scan trigger value section allows to define the reaction time. The default value of "0" is the most immediate since it sends an alarm at first connection attempt.
- The last section is to display a banner in case of intrusion attempt. It doesn't work with stealth modes.

When finished with that part, you're almost done. Again to avoid false alarms and huge logging, check the portsenentry.ignore file. You can add there your local network address with netmask bits, or the IP address of a few machines.

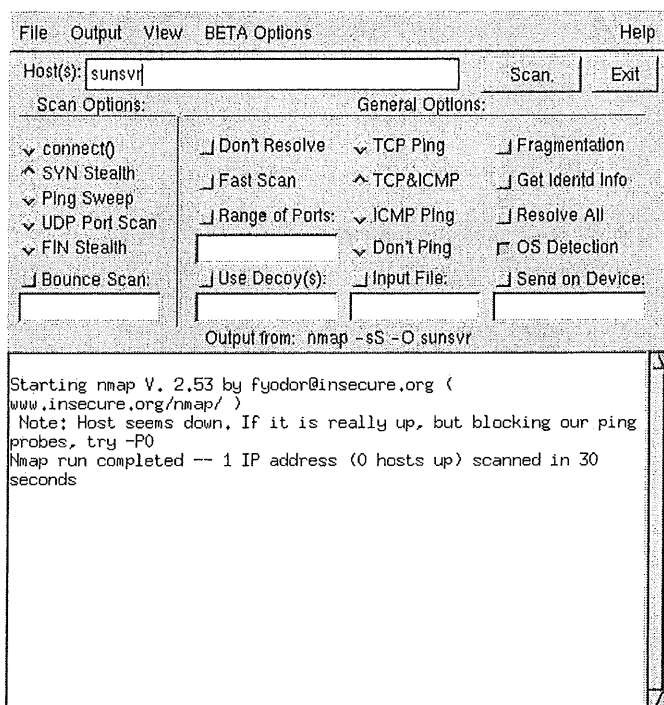
That's it ! You can now "make yoursystem" and "make install". The install script does the whole job since it gives restricted permissions to the directory and the files. Now, enough to run portsenentry with the options of your choice. Check the logs to see what happened. If everything went well, portsenentry has started and

has hardened the host security.

WHAT HAPPENS THEN ?

Using tools such as nmap, you may now simulate a port scan of your newly secured host. (More on nmap here:

<http://mercury.chem.pitt.edu/~tiho/LinuxFocus/English/July2001/article170.shtml>). This is the answer you can get :



That is, this machine doesn't give much information to the cracker !

Usually, this attempt should have provided you with the OS version and the open ports of the scanned host. This kind of information is quite useful for a cracker since he knows better where and how to attack. With portsentry running with a scan trigger value of 0 in advanced mode, blocking is immediate. Running with a value of 1 in normal mode, the attacker would have got a list of open ports and probably the OS type running on the machine. But then, trying to connect to a known open port would fail. Portsentry has blocked the port after first attempt. Simple but quite effective !

That is, if you only use TCPWrappers, the attacker IP address has been written into the /etc/hosts.deny file. If you decided to drop the route redirecting it to a dead host, portsentry is still efficient but doesn't fully prevent from UDP attacks.

If you use a packet filtering tool, the attacked host drops the packets coming from the attacker through this tool. Of course, this is the best way to benefit from the full power of portsentry. Once again, the advanced mode on Linux is the ideal one. This doesn't

mean portsentry is not as good on other OSes, it's just a bit less "performing"... well, not really !

To understand better, check the logs following different intrusion attempts to different OSes running portsentry in different modes. That says all !

According to the modes used, the logs will provide some false alarms. To improve the quality of the information, you'll have to work on the config file. After some testing, you should get exactly what you wish.

CAN YOU LIVE WITHOUT PORTSENTRY ?

Definitely NOT ! There are a lot of IDS, free or not. Most of them are quite good. Let's mention the well known snort, available from <http://www.snort.org>.

Usually, those IDSes are passive ones. That is, they are used to inform you about intrusion. That's what they have been made for. For instance, snort is quite sophisticated, since it relies on rules scripts, thus allowing you to write your own. Snort website also provides a rules database. And, by the way, if you aren't afraid of huge logs, you can even run both portsentry and snort. It isn't that silly !

However, the strength of portsentry is that it is an active IDS. The way it can react is really breathtaking. Combining it with a packet filtering tool, makes it the mandatory tool. We could ask the question : is portsentry the complement of packet filter or is packet filter the complement of portsentry ?

Anyway, there's a lot more we could say about portsentry. For instance, you must be very careful with advanced mode UDP scan detection. Do read the doc (again !) to understand the problems related to UDP. Solaris may cause some problems as well, since it uses very high ports in the 32700 range (either TCP or UDP). But, this article is only an overview and in no way a complete review of this great piece of software.

Computer security is a concern, and not only for sysadmin. The incredible growth of the Internet makes things worse. That is, the danger of intrusion is everyday bigger. The free software community provides us with great tools: portsentry is one of them. Don't leave them aside! But, remember, these tools only reduce the risk and in no way provides you with a 100% secure system.

This is another reason to investigate the security tools area. Last, but not least, people at Psionic are working on a commercial product and tell us "we won't be disappointed". Wait and see, then ! In the meantime, go and get portsentry 1.1, you won't be disappointed either.

People at Psionic do a great job. By the way, hostsentry is also worth testing... and using.

We are living in a great time, aren't we ?

REFERENCES

The following websites are absolute references when security is concerned. Furthermore, they can send you to some more websites, those last sending you to some more websites... You know, the Russian dolls? The only trouble if you want to read the whole stuff is that you won't live long time enough !!!

<http://www.linuxsecurity.com>
<http://www.sans.org>
<http://www.infosyssec.org>
<http://www.securityfocus.com>
<http://www.cs.purdue.edu/coast/hotlist/>

This article is re-printed with permission. The originals can be found at:

<http://mercury.chem.pitt.edu/~tiho/LinuxFocus/English/September2001/article214.shtml>

The State of Corporate IT: A case for Linux

Author: Paul D. Sullivan <paul.sullivan@anandtech.com>

We are in the midst of a severe corporate downturn, where the focus is more on the bottom line than perhaps ever before. As cost-cutters keep searching for ways to trim expenses, they often find themselves looking at one of the largest non-personnel related areas: IT Infrastructure.

Computer systems are such an integral part of most modern businesses that even the slightest hiccup can cause a serious financial hit. Most workers need these computers to do their jobs and the moment systems go down, productivity hits a wall. Not only do corporations have to be concerned with stability, they are becoming more and more concerned about the ongoing costs of running these systems.

The hardware is only the first step in creating this infrastructure. Once the hardware is in place, you need the software to get it all running in a productive manner. For nearly a decade, the momentum has been towards a Client/Server model based on the Microsoft Windows NT platform and for many, the move has meant an easier to administer, easier to implement configuration. But as NT has become more entrenched and companies have become more dependent on it, corporations have had to endure a drastic increase in overall costs.

Part of the increase has been hardware related. As NT has grown and become more capable, it has demanded much faster and more robust systems to work its magic. Another part of the increase has been security related. Corporations and consumers can be a very demanding bunch, and fairly or unfairly, Microsoft has been faced with the task of trying to

make NT all things to all people. In doing so, they have not been able to devote as many resources to security as they would have liked, and as a result, NT has proven to be less secure than originally hoped for. Their web server software in particular has been under assault from individual and organized attackers, and efforts have been increased to combat these intruders and to shore up the front line of defense.

But by many accounts, the largest cost of ownership increases that corporations have faced have been licensing related. As NT has become a mainstay, licensing terms have become more specific and more expensive. In addition, the explosive growth of the internet has brought security and reliability to the fore, and shoring up an evolving infrastructure can become prohibitively and increasingly expensive.

ONE TALE OF NT'S JOURNEY INTO CORPORATE AMERICA

To help make this evolution more understandable, we will use an example based upon the experiences of a corporation with a presence in Washington State. This company currently employs some 7,000 people at its primary site and had made the transition from a combination of Unix/Novell software to Windows NT. The move was cost justified based upon the ease of administration and a reduced cost of ownership, but years into the transition, administration and licensing costs soared and they were faced with some harsh realities, particularly when the market took a downturn and belt tightening became a necessity.

Initially, the company was approached by Microsoft and pitched on the idea of moving over to the NT Server platform. As any good company would do, Microsoft gave them a very hard sell and did an excellent job of convincing them of the potential benefits. One of the biggest parts of the pitch was the generous support that MS pledged to provide to corporate clients, and that support was perhaps the turning point in the decision to make the change. Direct support from IBM, Sun and Novell was becoming more expensive, even as it became harder to obtain. The company worked on a fresh agreement with Microsoft and took the plunge.

Originally, the licensing agreement called for a \$20 per seat annual fee for each client that would be accessing NT servers. Originally, two NT multiprocessor NT servers were put in place, each hosting Mail, Internet, File and Print requests. Costs for each server were based upon the number of clients accessing those servers at any one time. Since the company had their employees in separate shifts, only one half of their total employees would be capable of using those servers at any given time. After an initial analysis, it was determined that at no time did the number of concurrent users exceed 2,000. The drafted licensing agreement called for 2,000 concurrent licenses at \$20 per annum, for a total of \$40,000.

Over the next few years, as the transition from the old server software to the new became complete, changes to the licensing and service agreements were introduced. Uptime over those first years was not nearly as much as the service agreement had called for, and the cost to the company was becoming severe. When time came to renew the service agreement, the company sought more assurances and tighter uptime requirements. In response, Microsoft cited increasing demand on the servers and indicated that in order to ensure proper service and support, each server would have to be limited to one of three primary functions: Internet/Mail, File Serving and Print Serving. NT 4.0 performed best when each of those tasks was handled by a dedicated server and system integrity could only be guaranteed if the functions were separate.

In addition, the case was made that with the split of responsibility to multiple servers, the number of concurrent users could not be accurately determined and that it would be necessary to pay licensing fees for all of the clients that might use those servers. Employment at the company had increased by some 1,000 workers and network usage had increased along with it. When all was said and done, the company was asked to expand from two to six servers and to pay client access fees for a full 5,000 users on each of the three primary servers.

Under this proposal, annual licensing fees would increase from \$40,000 to \$300,000, but uptime performance would be guaranteed at a specific rate and there would be rebates should those rates not be met on a consistent basis. The company had invested over a million dollars to make the switch from the old to the new, and at this point, going back was not a viable option. Reluctantly the agreement was made and they moved forward.

The next couple of years saw a dramatic increase in data storage requirements and internet use as employment rose to nearly 7,000. The server redundancy helped ensure a higher level of uptime, but maintenance costs were going up as the internal IT team spent more hours working on the extra units that did go down, prepping them to go back up again. As redundant servers went in and out of service, data synchronization was becoming more critical and ensuring data integrity became an even costlier proposition.

During further licensing negotiations, Microsoft proposed that the company transition away from other suites and applications to Microsoft Office. In exchange for this move and the earlier commitment to the NT server line, Microsoft would give them a significant break on site licensing for these applications. They would even aid in transitioning their data warehouse from Oracle to SQL Server. At the time, the company took them up on the Office licensing bundle but skipped on the Oracle conversion. They would ride the market with the infrastructure they now had and do some long term evaluations before making any further commitments

to expanded licensing agreements.

THE TIDE SHIFTS...

As tends to happen with fast growing companies with demanding shareholders, there were some major changes in the ranks of upper management. With those changes came a more disciplined approach to cost structures and an increased scrutiny on exploding expenses such as IT budgets. It so happened that in the bi-annual review, one of the topics of discussion was the proposed transition to Windows 2000. Microsoft had proposed a very comprehensive package, but the cost conscious team leaders were hit with a fairly serious case of sticker shock. The decision was made to pull together the primary department heads, key IT staff and a team of Microsoft representatives to go over the possibilities.

After extensive meetings, some significant concerns began to surface. The proposed transition to Windows 2000 would be much more than a simple upgrade, but would actually constitute a paradigm shift in the way domains and assets were handled and managed. Active Directory (AD), a new and ambitious idea, was certainly not well established and at this point in the discussion, did not really seem ready for prime time. There appeared to be many unanswered questions on the MS side, and their "leap of faith" mantra was falling on skeptical ears.

According to the MS proposal, the transition to AD would eventually involve a complete shift of all internal systems to updated software, in part because Active Directory was not designed to be backward compatible with earlier Windows 9x client software. At the urging of MS, the company had earlier made a very substantial investment in the Primary Domain Controller (PDC) and Backup Domain Controller (BDC) paradigm established by earlier NT iterations. They had purchased the machines and the licenses, had passed dozens of their IT staff through the extensive training outlined by Microsoft and had literally worked years to develop a functioning asset management program utilizing IBM's Tivoli software. A change to the AD model would require another huge cash and resource infusion, and might possibly negate many of the benefits realized from the earlier investments in the PDC/BDC model.

As if all of this was not "Red Flag" enough, there would be some serious changes in the licensing agreement that would now cover all types of external and internal network accesses including terminal services, remote dial-in and the use of Virtual Private Networks (VPN's). Further, another proposed modification would take into account the number of accesses and transactions conducted over servers using MS hosting software. When it was all put together, the cost implications were absolutely staggering.

SINK, SWIM OR TREAD WATER?

The company in this example was at a critical point in the decision making process. They had made the choice some years back to migrate to the NT platform and invested heavily in the infrastructure. But now they were forcing themselves to take a long, hard look at their decision before they would commit any further. Cost analysis had shown that even though the company had poured a great deal of money into software, hardware and institutional retraining, the benefits were not nearly as apparent as promised.

The fact that the NT platform had not proved itself to be as scalable or reliable as promoted was a major thorn in their side. The redundancy suggested by Microsoft had helped increase uptime, but increased the already high maintenance of the growing server clusters. While the previous Unix and Novell platforms had handled file, print and mail servers on a single server, NT now needed one machine for each service plus a dedicated backup for each. Hardware costs were not the real concern – it was the licensing and maintenance requirements that hit the hardest.

The modern workforce was changing, and remote access was becoming much more important for collaborative efforts and virtual workstations. The excessive costs of purchasing licenses for each and every terminal server, remote dial-in and VPN access was simply not going to be possible with the start of the downturn in the tech sector. Though E-commerce was not as explosive as predicted, the web was becoming a very significant tool for internal and external customers alike, not to mention vendors and other third parties. Paying fees for each of these transactions would severely blunt the effectiveness of the entire process, actually making it more cost effective to take a step back and do things the old way.

The company was knee-deep in the mire here, and had to face some tough decisions. Luckily, they had some options. Linux had been gaining a steady groundswell of support over the past few years and had some serious advantages in terms of cost to benefit ratios. Not only was the software free, but it would run on existing hardware and could actually be tweaked and recompiled to maximize performance in key areas such as file and print serving. Plans to advance into a modified agreement with Microsoft were to be put on hold in favor of further exploration into Linux. It would be a decision that they would not regret.

THE ADVANTAGES OF LINUX

Red Hat was key in helping them realize the benefits to the bottom line. Within 60 days of the first overtures, they were on site with a demonstration that completely blew the corporate team away. Red Hat brought a single Pentium class system for a site visit and thanks to the early legwork their engineers had done, were able to integrate the box into the network and take over all file and print server requests for one busy segment within four hours. The system ran for the next 10 business days without any downtime,

something NT machines had not been able to do very often. All issues that did come up were fixed on the spot without a single kernel restart. File and print transactions were stored in ques and processed without incident. Samba allowed the Linux box to seamlessly integrate into the file network and actually increased overall performance. Nightly backups were performed from the master NT server without any sign of incompatibility. Print jobs were also handled seamlessly with fewer delays and error messages along the way. This limited demonstration was an absolute success and had most of the corporate advance team nodding their heads in approval.

Additional demonstrations followed that focused on developers and system administrators. Using a Linux system with a combination of existing products and a newer application called VMware, developers were able to write, compile and test code on a variety of distinct platforms from a single machine. When code did crash, it affected only one environment, which could quickly be killed and restarted without interfering with other processes. Since the host GUI was not tied to the Linux kernel, but was instead run as a separate process, even the most complex and sensitive graphical development would not bring down the machine. In each case, individual processes were simply killed and restarted to a fresh state. Productivity benefits were obvious.

By activating an Apache web server on the same machine, development, administration and testing of the corporate intranet could be handled in real-time. Quick fixes could be made in code windows, loaded in the appropriate Apache folders and tested on multiple platforms with multiple browsers within minutes. This level of stability and flexibility was something these developers had been craving for some time. Part of the excitement involved the bottom line as well. With the exception of VMware, all of the Linux software was essentially free and did not have any of the restrictions imposed under NT EULA's. Developers were able to clone and distribute development environments to other machines across the hall or across the country without fear of violating licensing terms.

In fact, with Linux, there would be no licensing fees at all. The proposed file servers, print servers and web servers to be hosted under Linux did not require the purchase of any client licenses whatsoever. There would be no E-commerce transaction fees, no distribution limitations and no expensive application bundling requirements. Developers could use a variety of existing applications, from Star Office to SQL databases to C++ IDE's, all without additional fees. Contractors could be provided with all the tools they needed with no additional impact on cost.

During these developer and administration demonstrations, it was found that users already familiar with Unix seemed to feel right at home in Linux, so extensive retraining would not be needed. The KDE environment was similar enough to Windows so that even novice users got the hang of it

pretty soon. Since Star Office was also free for use on Windows machines, entire project teams could be assembled and provided with some of the same key tools, regardless of platform - all without impacting the bottom line. The possibilities were becoming obvious, as were the advantages. Linux provided more freedom and more flexibility at a lower initial and long-term cost. Linux gave them an alternative - one that could be explored and researched without the need to allocate additional capital funds. It was a pretty easy decision to continue exploring this new platform.

LINUX MAKES AN IMPACT

In addition to all the benefits and possibilities mentioned earlier, Linux gave this company a bargaining chip in license negotiations with Microsoft. Earlier, they had passed on the Windows 2000 and Office 2000 upgrades, waiting instead to see what the future held after the evaluation. When the move to XP was being touted by Microsoft during subsequent meetings, they found that they were faced with some surprises. The proposed fee structure was radically different and established what amounted to a cost penalty for those who chose to stand pat instead of upgrading their operating systems and office suites to XP versions upon their initial release.

Currently, they had a large mix of Windows 9x and NT 4.x clients successfully running Office 97 software on laptops and desktops. They had stayed on top of bug fixes and system patches and found that for the most part, their existing infrastructure performed well enough as it was. They had maintained their NT 4.x Server infrastructure as well, opting not to move to Active Directory.

During this difficult time, Red Hat had proven to be a helpful ally. Instead of trying to push a whole-scale replacement of the infrastructure, they had worked to supplement it. Over time Linux brought more security, improved load balancing and an overall reduction in the growth rate of IT spending. Point of sale terminals were reliable, easy to manage and did not incur additional transaction costs. Their remote access and VPN configurations handled an ever increasing load with a higher degree of reliability and a lower cost. Their intranet had been transitioned over to Linux, and as a result cost less to maintain. It also eliminated interference with IIS based consumer and vendor systems accessed from outside of the company.

Through a series of such modifications, they had been able to establish and maintain a more stable, more cost effective configuration. Their network was more flexible and more able to meet the needs of a changing marketplace. Projects could be isolated to their own LAN or WAN segment without impacting other services and teams of experienced Unix/Linux workers could be called upon when NT resources were scarce. As a result, overall TTM (Time To Market) was reduced for mission critical consumer

applications and customer satisfaction actually rose in the midst of explosive growth.

Linux was not the right tool for every job, but it certainly had proved its mettle as a cost effective alternative and helped give them some breathing room as they worked to bring soaring IT costs under control and reduce TCO (Total Cost of Ownership). It was ironic that only by turning to an alternative operating system were they able to realize some of the cost savings promised them when they initially switched over to NT. Linux had not only given them tangible benefits, it had increased confidence in their ability to manage their own systems.

This was important because over time there had been a growing fissure between what Microsoft had originally promised and the proposals they were making today. The constant tinkering with licensing agreements, the perpetually increasing fees and the imposition of bundling and usage restrictions had generated a lot of bad-will. The Microsoft of old that had come knocking on their door with friendly overtures was no more. In its place was a company that stifled their clients with ever-increasing pressure to upgrade or face the prospect of paying higher fees and receiving reduced levels of support.

So when the time came to make a decision on the transition to XP, they felt they were in a much stronger position. They had found a willing partner in Red Hat, a viable alternative in Linux and a sense of control over their own infrastructure that had previously been lacking. Though they might face higher licensing costs later on, they opted to again bypass the proposed Microsoft solution in favor of standing pat.

SUMMING IT ALL UP

Some months later, with the market still soft and the bottom line increasingly important to shareholders, the team feels they made the right decision. The proposed licensing agreements would have required a complete switch to new versions of the Microsoft XP software, increasing deployment costs. The system requirements of Office and Windows XP would have mandated a substantial investment in the purchase and deployment of new hardware and the transfer of system data. Changes to the system interface would have required worker retraining. Changes to the core OS would have rendered a variety of third party software and utilities unusable, dramatically impacting productivity and further increasing upgrade costs.

In addition, Microsoft was being put under intense pressure from Federal Regulators, the courts and consumers. They were at odds with Sun over Java, with vendors of the status of icons and with consumers over security and product activation. Hackers were finding holes in their server big enough to drive a truck through, and security experts like Steve Gibson were bringing other serious flaws to light. They questioned the need to include Direct X and other consumer-friendly services into the server

version of their OS. They questioned the extensive integration of IE 6 and Smart Tags. They questioned the usage requirements of their Passport software, a key part of their XP and .Net strategies.

Sometimes, when there's smoke, there's fire and heading in another direction can be the smart thing to do from a client standpoint. Companies have been dropping like flies in this marketplace and only the nimble seem to be able to survive. You have to be willing to challenge existing paradigms.

Change does not always have to be a frightening thing, and it is always a good idea to have alternatives at your disposal in the corporate world. Having "all your eggs in one basket" has been considered a risky proposition for a lot longer than computers have been around, but the principle is still sound and timely. In the case of the company we used as a basis for this example, thinking outside the box paid off.

As a result of their willingness to look beyond, they now have a more cost-effective, more stable and more predictable infrastructure in place. They have been able to benefit from the hard work of the Linux community and the support of companies like Red Hat. They have been able to establish and maintain key relationships with forward thinking companies like Dell, who started bundling and supporting Linux on their server machines early on.

Because they were willing to open their eyes to new ideas and challenge convention, they have been able to hang on where others have not. If other companies facing growing infrastructure costs are to survive these difficult times, it may be a good idea for them to do some evaluations of their own. After all, Linux is a free download and Red Hat is only a phone call away...

This article is re-printed with permission. The originals can be found at:

<http://www.anandtech.com/showdoc.html?i=1527>

Installing Debian on 4 MB Ram System

Author: Jens Scheidtmann <JensScheidtmann@web.de>

Version 0.1

Dear readers,

recently I reactivated an old B/W notebook with 4 MB Ram lying around at my home: I bootstrapped Debian on it. The bootstrapping process was a bit more complicated as a normal install of Debian. This miniHOWTO outlines how I have done it. As this bootstrapping has happened some time ago and I am recalling it from memory there is as usual no guarantee that it will work. If you find any errors (be it typos, grammar or wrong facts) please tell me. I am currently subscribed neither to debian-user nor to debian-devel, so please cc your replies to

JensScheidtmann@web.de (It would overwhelm my one and only mail account -- I am currently waiting for a mail account with more space). This document is copyright (C) 2001 Jens Scheidtmann.

PREREQUISITES

- A normal computer with some free disk space. Not only free space in a partition, but an area of your hard disk without a partition. It should run Debian of course.
- A notebook or computer with ≥ 4 MB RAM (and < 8 MB) with an HD of ≥ 300 MB, a parallel port and a floppy drive.
- A LapLink cable (at least this howto assumes it, but you sure can use other means to connect the two computers).
- An Internet connection or a Debian Install CD Set.
* A version of the smalllinux boot/root disks (search freshmeat.net).

OUTLINE OF THE PROCESS:

- 1) First you will use your normal computer and prepare a tarball of the Debian base system.
- 2) Then you will boot the notebook with boot/root disks, partition its hard disks and transfer the prepared tarball to it.
- 3) Make it bootable from the notebooks HD.
- 4) Start using Debian on your notebook.

Note: Before following the process described here, you can try to install

Debian from the installation disk set. Make sure that you replace the kernel on the boot disk with a kernel which does not consume so much memory. I don't know exactly where I ran into problems (because I did this install described in here a while ago), but I tried hard and fell back on the approach described in this howto, so good look and please tell me if you get it working.

PHASE I

Prepare the Tarball.

- 0) Boot into smalllinux on your Notebook.
- 1) Look at your notebook and find out how big the HD is using fdisk. Decide how many space you will reserve for swapping and how many space is left for the linux partition. Use plenty of swap space! Mine is 35 MB. You should make it at least 20 MB big, because this partition will first contain a copy of the tarball we prepared on the normal computer. So partition the disk according to your choices. Format BOTH partitions as ext2.
- 2) Create a partition with a comparable size as the final linux partition on your normal computer. The exact size does not matter.
- 3) Install the Debian base system into this partition using the install floppies, a CD of your installation CD set or use whatever is your preferred way to install Debian.

- 4) Boot into your normal Linux System and mount the partition. Make a nice tarball out of it (as root), while preserving permissions:

```
# cd /mnt/point; tar cvpzf /tmp/ball.tar .
```

PHASE II

Once you got the tarball, you can proceed by making a network connection between the two computers.

- 0) Take a break.
- 1) Connect the two computers with your LapLink Cable.
- 2) Make sure the kernel of your normal computer has got the "plip" module available or compiled in.
- 3) Edit /etc/hosts and add two lines for your normal computer and your notebook:

```
--/--
192.168.0.1   father
192.169.0.2   baby
--/--
```

- 4) Fire the connection up on your normal computer:


```
# ifconfig plip0 father pointpoint baby up
```

 If you are still using a 2.0.X kernel use plip1 instead. You will then have to add a route, too: "route add baby dev plip1".
- 5) Similarly fire the connection up using smalllinux from the notebook. Here you will have to use "plip1", raw IP addresses and add a route. You should now be able to send pings between the two computers.
- 6) mount the supposed swap partition as ext2 and open an ftp connection to your normal computer. Download ball.tar into the swap partition.
- 7) mount the final linux partition and extract the tarball into it:


```
# cd /mnt/point/final; tar xvzpf /where/ever/it/is/ball.tar
```
- 8) umount the swap partition and do a mkswap with it.

PHASE III

Now the Debian base system is installed on the laptop, you now have to make it bootable.

- 0) Prepare a bootfloppy for you notebook by dd-ing the smalllinux kernel onto a floppy disk. Set the correct boot device (the final partition) with rdev:

```
# dd if=smalllinux.kernel of=/dev/fd0
# rdev /dev/fd0 /dev/correct.partition
```

- 1) Start the notebook with it and keep your fingers crossed that everything went alright and you don't get a kernel panic.
- 2) Copy the kernel to your harddisk, edit lilo.conf and run lilo.
- 3) Check whether the notebook starts without problems.

PHASE IV

Use the notebook. Here are some tips:

- Compile a lightweight 2.0.x kernel on your normal computer using kernel-tools. If it takes more than 500 k it is not lightweight. "dpkg -i" it on the

notebook. (fdisk -l will not work, but hey, who cares.)

- ftp /etc/apt/sources.list and /var/state/apt/lists/* from your normal computer. Read /usr/share/doc/apt/offline.txt.gz and use the wget method, but ftp the packages directly to /var/cache/apt/archives. Try to avoid installing, because apt-get install swaps to death.
- zgv is a svgalib picture viewer, and bmv is a svgalib ghostscript frontend.
- emacs20 works but takes ages to start.
- I even installed tetex.
- There is a "tiny X" out somewhere, but I haven't tried it yet and I don't know when I will.

You may freely distribute and alter this document under the GNU Public License 2 or above

This article is re-printed with permission. The originals can be found at:

<http://7thguard.net/files/DebianHOWTO.txt>

wxPython for newbies

Author: Michael Roberts <michael@vivtek.com>

ABSTRACT

You can write a Python script in minutes and have incredibly nice-looking GUI apps for your desktop. This article shows you how to use one Python-savvy GUI library, wxPython, to do just that. Impress your friends and neighbors!

INTRODUCTION

This article is about wxPython, but wxPython is really the combination of two things: the Python scripting language and the wxWindows library of GUI functionality (for an intro to wxWindows, see "Looking through wxWindows", here on developerWorks. The wxWindows library is a C/C++ library that abstracts GUI functionality for maximum portability. So wxWindows apps can run natively under Windows, UNIX with X, KDE or Gnome, or whatever wxWindows has been ported to (unfortunately that doesn't yet include the Macintosh). Of course Python, as a scripting engine, is also extremely portable (and does run on the Macintosh, not that it helps if you want to write desktop GUI code). Combining the two means that your wxPython apps are not only quick and easy to write, they'll also run without changes on either Windows or UNIX.

You might be thinking, "But that's why I have Java. Java is also portable." Well, if you've ever tried to install a Java application on Windows, you might recognize that this isn't entirely the case. The Java virtual machine is big, it doesn't always work the way

you want, and worst of all, Java windows aren't real windows so interaction with the host system is always at arm's length, so to speak.

Python, on the other hand, has a relatively small footprint. The wxPython library's windows are real live native windows and can do anything native windows can do, allowing your wxPython program to be a much better-behaved citizen. And the whole shebang can be wrapped into an easy installation package. Maybe I'm just a curmudgeon, but I find all this to be much easier than the same thing under Java.

But you probably haven't heard about Python on the desktop, although it's catching on in the server-side programming community as the new kid on the scripting language block (especially in connection with application server frameworks like Zope). People are really getting caught up in Python mania. And the nice thing about Python is that unlike other scripting languages, it's an object-oriented language from the ground up. So you're not losing that tasty OO goodness by ignoring Java in favor of Python.

The world's tiniest wxPython program, dissected! So it sounds cool so far, doesn't it? But let's look at some code and you'll see what I mean.

I'm inserting some line labels in this sample to make it easier to talk about. They're not part of the code; that's why they're in square-brackets: [].


Listing 1. A tiny code sample

```
[1] import sys, os
[2] from wxPython.wx import *
[3] class main_window(wxFrame):
[4]     def __init__(self, parent, id, title):
[5]         wxFrame.__init__(self, parent, -1, title,
size = (200, 100),
style=wxDEFAULT_FRAME_STYLE|wxNO_FULL_REPAINT_ON_R
ESIZE)
[6]         self.control = wxTextCtrl(self, -1,
style=wxTE_MULTILINE)
[7]         self.Show(true)
[8] class App(wxApp):
    def OnInit(self):
        frame = main_window(None, -1,
"wxPython: (A Demonstration)")
        self.SetTopWindow(frame)
        return true
[9] app = App(0)
    app.MainLoop()
```

Looking at this line-by-line, it's easy to see how it works. That's why it's the world's tiniest wxPython program (I've taken out all the extraneous details). All this code does is create a window with an edit field in it. You can edit in the field, but obviously there's no way to save things. Here's what the application looks like under Windows (I've typed a little into it):

SO LET'S LOOK AT THE CODE

Lines 1 and 2 are important, and we'll need them later. Actually, we don't need sys and os in this example, but since they get used in nearly every program, I tend to include them first and ask



questions later. The more interesting part is line 2, where everything from the wx library of wxPython is imported. The wx library is, of course, where the basic wxPython declarations are, including basic classes like frames and applications.

Notice that there are no trailing semicolons on these lines. If you write any Perl at all, it will take you a little effort to get into Python typing habits. But wait, it gets stranger.

In line 3, we define our first Python class, main_window. The main_window class is derived from the wxFrame class that was defined in the wx module. As you can guess, any window is a wxFrame class.

In line 4, we define the initialization method, the only method for main_window that needs to be defined (everything else is in the wxFrame class, of course). The initialization method takes parameters parent, id, title, and of course self as a reference to the object. The self parameter is the first parameter of all Python methods.

By now, if you write C/C++ or Perl code at all, you will be starting to wonder why there are no curly braces anywhere. Oddly enough, Python treats indentation as syntactically significant. Any statement that has sub-statements ends in a colon, and all indented lines under that line belong to the colon-terminated statement. When the indentation reverts to the original indentation, the block is finished. Even more strangely, this arrangement actually seems to work, and beginning programmers find it easier to use than curly braces to indicate structure. Fortunately, it also means fewer keystrokes, so advanced programmers can get used to it too.

So the indented lines in 5, 6, and 7 belong to the method defined in line 4. Respectively, they call the initialization routine of the wxFrame parent class (doing all the hard work of actually setting up a window), define a control to fit on the new main window, and make sure the window is visible.

Unless specifically told not to, wxPython controls resize with their parent window. If you've ever written any Microsoft Windows code that resizes controls to match the parent window, be it in C/C++ or in Visual Basic, then you will immediately learn to love wxPython. The practical meaning of this is that line 6 implements an editor every bit as capable as Notepad. OK, that's not all that impressive. But it would still take just a few more lines of C/C++ code to do that, now wouldn't it? Remember -- this isn't just a fixed-size dialog we're doing here, it's a real Windows-y application with the editor in the body of the window

and perfectly resized.

So let's get on with it. Line 8 defines the App object, which derives from the wxApp class, of course. It specifies that the application object, when started, creates a main_window object and sets it as the top window.

Finally, now that our classes are defined, we create the application object and start its main loop running. If you've done any C/C++ Windows programming, you'll recognize the MainLoop method as the normal event loop of all Windows programs. The style of this startup code is definitely in keeping with a scripting language. Remember that the Python interpreter reads this file line by line and executes as it goes. So once the classes are defined, we simply call them in the script.

And that's it. In fifteen lines of code, we have implemented a simple text editor that will run without modification on Windows or UNIX. And it's extremely easy to add more features, which is good, because that's what we're going to do in the next section.

WHAT ABOUT PERFORMANCE?

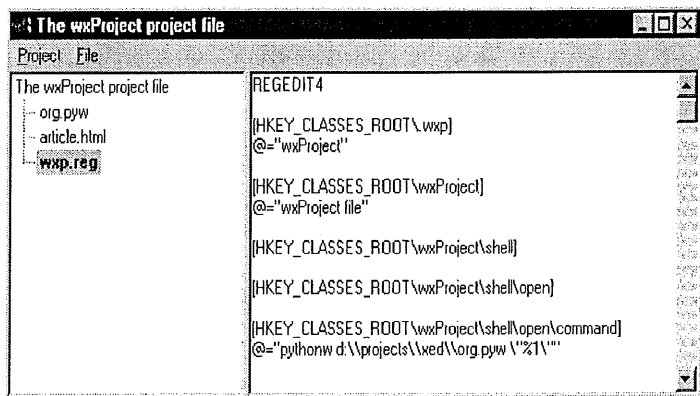
You may be thinking, "But an interpreted language is so inefficient. Any sizable program would be too sluggish." And that's correct, to a certain extent. In fact, any performance-impacting code is usually re-implemented in C/C++ and linked into the Python interpreter, which is extremely easy to do. So Python is generally used as the glue to bind functional modules with GUI presentations (or with application server functions, if you're into that scene.) But as glue, Python is extraordinarily efficient. You can implement serious programs in very little time with Python, and they generally stay fairly readable weeks later due to the object orientation and the limited opportunity for creativity in formatting.

Moreover, if you should doubt the idea of using an interpreted language to tie a large program together, you might consider that this is exactly how early versions of Microsoft Word were implemented. At least as recently as version 6.0 for Windows, the Word Basic functions were actually the only compiled code; the GUI was built with a pcode-interpreted language. (Effectively, MS Word 6.0 was written in Word Basic, which is one of the reasons it was one of the earliest desktop programs with a built-in interpreter -- pretty slick design.)

What you lose here in performance (in practice, very little), you easily regain in ease of implementation and (more importantly) ease of customization. In fact, by including Python you have already automatically included a scripting language that you can easily expose to your users, and which has been demonstrated to be easy to learn for newcomers to programming. If you leverage this, you can end up creating world-class software with a lot less effort than you're expending now. If it worked for Microsoft, it can work for you!

SOMETHING MORE INTERESTING: A RUDIMENTARY PROJECT ORGANIZER

Anyway, time to put my money where my mouth is and show you a program that actually does something worthwhile. This program allows you to create a list of text files you call a project. You can edit the files and save them. More importantly, you can easily see how the basic organizer could be enhanced further. I use an extended version as a front end for CVS (the standard open-source version control system). Here's the code below. Our (still rather small) application has gone from 15 lines of code to around 300, but it does a lot more now too.



Listing 2. The rudimentary project organizer

```
#!/bin/python
import sys, os
from wxPython.wx import *
# Process the command line. Not much to do,
# just get the name of the project file if it's given. Simple.
project = 'Unnamed'
if len(sys.argv) > 1:
    project = sys.argv[1]
def MsgBox(window, string):
    dlg = wxMessageDialog(window, string, 'wxProject', wxOK)
    dlg.ShowModal()
    dlg.Destroy()
class main_window(wxFrame):
    def __init__(self, parent, id, title):
        wxFrame.__init__(self, parent, -1, title, size = (500, 500))
style = wxDEFAULT_FRAME_STYLE | wxNO_FULL_REPAINT_ON_RESIZE
#
# Set up menu bar for the program.
#
self.mainmenu = wxMenuBar() # Create menu bar.
mainwindow = self
menu = wxMenu() # Make a menu (will be
the Project menu)
exitID = wxNewId() # Make a new ID for a
menu entry
menu.Append(exitID, 'Open project') # Name the ID by adding it
to the menu.
EVT_MENU(self, exitID, self.OnProjectOpen) # Create and assign a
menu event
exitID = wxNewId()
menu.Append(exitID, 'New project')
EVT_MENU(self, exitID, self.OnProjectNew)
exitID = wxNewId()
menu.Append(exitID, 'E', 'Exit program')
EVT_MENU(self, exitID, self.OnProjectExit)
self.mainmenu.Append(menu, 'Project') # Add the project menu to the
menu bar.
menu = wxMenu() # Make a menu (will be
the File menu)
exitID = wxNewId()
menu.Append(exitID, 'Add file to project')
EVT_MENU(self, exitID, self.OnFileAdd)
exitID = wxNewId()
menu.Append(exitID, 'Remove file from project')
EVT_MENU(self, exitID, self.OnFileRemove)
exitID = wxNewId()
menu.Append(exitID, 'Open file for editing')
EVT_MENU(self, exitID, self.OnFileOpen)
exitID = wxNewId()
menu.Append(exitID, 'Save file')
EVT_MENU(self, exitID, self.OnFileSave)
self.mainmenu.Append(menu, 'File') # Add the file menu to the
menu bar.
self.SetMenuBar(self.mainmenu) # Attach the menu bar
to the window.
#
# Create the splitter window.
#
```

```

splitter = wxSplitterWindow(self, -1, style=wxNO_3D|wxSP_3D)
splitter.SetMinimumPaneSize(1)
#
# Create the tree on the left.
#
tID = wxNewId()
self.tree = wxTreeCtrl(splitter, tID, style=wxTR_HAS_BUTTONS |
                        wxTR_EDIT_LABELS |
                        wxTR_HAS_VARIABLE_ROW_HEIGHT)
EVT_TREE_BEGIN_LABEL_EDIT(self, tree, tID, self.OnTreeLabelEdit)
EVT_TREE_END_LABEL_EDIT(self, tree, tID, self.OnTreeLabelEditEnd)
EVT_TREE_ITEM_ACTIVATED(self, tree, tID, self.OnTreeItemActivated)
#
# Create the editor on the right.
#
self.editor = wxTextCtrl(splitter, -1, style=wxTE_MULTILINE)
self.editor.Enable(0)
#
# Install the tree and the editor.
#
splitter.SplitVertically(self.tree, self.editor)
splitter.SetSashPosition(180, true)
self.Show(true)
# Some global state variables.
self.projectdirty = false
#
# Some nice little handlers.
#
def project_open(self, project_file):
    try:
        input = open(project_file, 'r')
        self.tree.DeleteAllItems()
        self.project_file = project_file
        name = replace(input.readline(), "\n", "")
        self.SetTitle(name)
        self.root = self.tree.AddRoot(name)
        self.activeitem = self.root
        for line in input.readlines():
            self.tree.AppendItem(self.root, replace(line, "\n", ""))
        input.close()
        self.tree.Expand(self.root)
        self.editor.Clear()
        self.editor.Enable(false)
        self.projectdirty = false
    except IOError:
        pass
def project_save(self):
    try:
        output = open(self.project_file, 'w')
        output.write(self.tree.GetItemText(self.root) + "\n")
        count = self.tree.GetChildrenCount(self.root)
        iter = 0
        child = None
        for i in range(count):
            if i == 0:
                (child, iter) = self.tree.GetFirstChild(self.root, iter)
            else:
                (child, iter) = self.tree.GetNextChild(self.root, iter)
            output.write(self.tree.GetItemText(child) + "\n")
        output.close()
        self.projectdirty = false
    except IOError:
        dlg_m = wxMessageDialog(self, 'There was an error saving the
project file.', 'Error!', wxOK)
        dlg_m.ShowModal()
        dlg_m.Destroy()
#
# Event handlers from here on out.
#
def OnProjectOpen(self, event):
    open_it = true
    if self.projectdirty:
        dlg = wxMessageDialog(self, 'The project has been changed. Save?',
'wxProject', wxYES_NO | wxCANCEL)
        result = dlg.ShowModal()
        if result == wxID_YES:
            self.project_save()
        if result == wxID_CANCEL:
            open_it = false
        dlg.Destroy()
    if open_it:
        dlg = wxFileDialog(self, "Choose a project to open", "", "",
** .wxp", wxOPEN)
        if dlg.ShowModal() == wxID_OK:
            self.project_open(dlg.GetPath())
            dlg.Destroy()
def OnProjectNew(self, event):
    open_it = true
    if self.projectdirty:
        dlg = wxMessageDialog(self, 'The project has been changed. Save?',
'wxProject', wxYES_NO | wxCANCEL)
        result = dlg.ShowModal()
        if result == wxID_YES:
            self.project_save()
        if result == wxID_CANCEL:
            open_it = false
        dlg.Destroy()
    if open_it:
        dlg = wxTextEntryDialog(self, "Name for new project:", "New
Project", "New project", wxOK | wxCANCEL)
        if dlg.ShowModal() == wxID_OK:
            newproj = dlg.GetValue()
            dlg.Destroy()
            dlg = wxFileDialog(self, "Place to store new project", "",
"", ".wxp",
wxSAVE)
            if dlg.ShowModal() == wxID_OK:
                try:
                    proj = open(dlg.GetPath(), 'w')
                    proj.write(newproj + "\n")
                    proj.close()
                    self.project_open(dlg.GetPath())
                except IOError:
                    dlg_m = wxMessageDialog(self, 'There was an error
saving the new
project file.', 'Error!', wxOK)
                    dlg_m.ShowModal()
                    dlg_m.Destroy()
            def OnProjectExit(self, event):
close = true
if self.projectdirty:
    dlg = wxMessageDialog(self, 'The project has been changed. Save?',
'wxProject', wxYES_NO | wxCANCEL)
    result = dlg.ShowModal()
    if result == wxID_YES:
        self.project_save()
    if result == wxID_CANCEL:
        close = false
    dlg.Destroy()
    if close:
        self.close()
def OnFileAdd(self, event):
    dlg = wxFileDialog(self, "Choose a file to add", "", "",
wxOPEN)
    if dlg.ShowModal() == wxID_OK:
        path = os.path.split(dlg.GetPath())
        self.tree.AppendItem(self.root, path[1])
        self.tree.Expand(self.root)
        self.project_save()
def OnFileRemove(self, event):
    item = self.tree.GetSelection()
    if item != self.root:
        self.tree.Delete(item)
        self.project_save()
def OnFileOpen(self, event):
    item = self.tree.GetSelection()
    def OnFileSave(self, event):
        if self.activeitem != self.root:
            self.editor.SaveFile(self.tree.GetItemText(self.activeitem))
        def OnTreeLabelEdit(self, event):
            item = event.GetItem()
            if item != self.root:
                event.Veto()
        def OnTreeLabelEditEnd(self, event):
            self.projectdirty = true
        def OnTreeItemActivated(self, event):
            go_ahead = true
            if self.activeitem != self.root:
                if self.editor.IsModified():
                    dlg = wxMessageDialog(self, 'The edited file has changed. Save
it?',
'wxProject', wxYES_NO | wxCANCEL)
                    result = dlg.ShowModal()
                    if result == wxID_YES:
                        self.editor.SaveFile(self.tree.GetItemText
(self.activeitem))
                        if result == wxID_CANCEL:
                            go_ahead = false
                            dlg.Destroy()
                        if go_ahead:
                            self.tree.SetItemBold(self.activeitem, 0)
                    if go_ahead:
                        item = event.GetItem()
                        self.activeitem = item
                        if item != self.root:
                            self.tree.SetItemBold(item, 1)
                            self.editor.Enable(1)
                            self.editor.LoadFile(self.tree.GetItemText(item))
                            self.editor.SetInsertionPoint(0)
                            self.editor.SetFocus()
                        else:
                            self.editor.Clear()
                            self.editor.Enable(0)
class App(wxApp):
    def OnInit(self):
        frame = wxWindow(Home, -1, "wxProject - " + projfile)
        self.SetTopWindow(frame)
        if projfile != 'Unnamed':
            frame.project_open(projfile)
        return true
app = App(0)
app.MainLoop()

```

FUN STUFF IN THE PROGRAM

I didn't number lines in Listing 2, because you're obviously here to read code anyway. I'll just generally talk about what the program is doing and what features of Python and wxPython it's using to do it. For a much more detailed exposition of Python, head to your local bookstore and pick up the O'Reilly book by Mark Lutz, or just read the doc pages which ship with Python (written by Guido Rossum, the actual author of Python).

The first interesting thing is that this app processes the command line. The list `sys.argv` is the command line, and you need to understand Python list syntax to work with it. The example shown is pretty basic, but should get you started.

Next, we've defined a `MsgBox` function just for ease of use when debugging. Note that the syntax for function declaration and class method declaration is precisely identical. The only difference is that a class method is required to take a parameter `self` that contains a reference to the object being called. (You don't have to call it "self", of course. But it'll be confusing if you don't.)

The fun really starts in the much more complicated `__init__` method. Here, instead of making a rinky-dink little control, we have to build a menu bar, attach menu events to callback routines, and build a

window splitter bar, tree control, and the edit control. You can skim over that and see how it all works. (If you haven't done GUI programming before, it will be kind of tough going.) Visual Basic does all this in a kind of behind-the-scenes manner, but you can still see it if you simply open a form file up in a text editor. If you've ever done GUI work in C/C++, this will all look eerily familiar.

Once we have the construction of the window out of the way, we get down to the actual code for the program. First off, after the `__init__` method are two methods used to load and save project files. There you can see how Python handles file I/O, with open and so forth. Note that actually closing files is a kind of relaxed thing -- as it happens, a file handle is simply a memory-managed object to Python and objects are referenced by count. As soon as a reference becomes invalid, Python knows it and cleans the object up, and the file gets closed automatically. There are cases where this can't be entirely trusted, so for instance where you're going to reopen a written file for reading, it's a good idea to close it explicitly. However, it illustrates the whole garbage-collection thing that non-C people like to talk about so much.

Another interesting feature of the load/save functions is that they catch exceptions generated by bad file calls (the IOError exceptions). I'll let that code speak for itself, but that's how you do that, boys and girls.

The remainder of the new code is event handlers for things that happen during program execution. And again I'll let the code speak mostly for itself. Note, however, the use of common dialogs (`wxMessageDialog`, `wxFileDialog`, and `wxTextEntryDialog`) to handle a lot of routine user interaction. These calls are a little different from the corresponding uses in "regular" Windows programming, so I'll just give you a couple of clues: first, the dialog is an object that is created by the appropriate call, displayed using `ShowModal()`, and then destroyed after use; second, the button that the user clicks is returned as the return value from `ShowModal()`, and other values are accessible using dialog-dependent methods. So for instance, the path selected by the user during the file dialog interaction is available as `dlg.GetPath()`. And that's rather different from the way the Windows APIs work. Just so you know.

THE DOCUMENTATION SITUATION

Documentation of the wxWindows API is ... under development, let's say. Unfortunately, the lack of good documentation is the greatest weakness of this platform, and I'm getting active in rectifying that situation. In another six months things will be much better. In the meantime, Python itself has some good documentation books available, and Python's online documentation isn't bad. When it comes to wxWindows, the C++ documentation is pretty much all there is. Sometimes it's rather mysterious how to make that make sense in the Python framework.

There are some Python-specific annotations in there, but in many cases you'll find yourself having to read the demo code that comes with wxPython, or simply having to ask questions on the mailing list. Fortunately, the people on the mailing list are "clueful" individuals.

And once you're past the initial learning curve, it's much easier stuff to work with than the same tasks in the Windows API. Believe me. This stuff is good.

RESOURCES

- For an intro to wxWindows, see "Looking through wxWindows" here on *developerWorks*.
- The wxPython home page links to the documentation, downloads, and related sites.
- This site is mirrored at ourceForge, and it has links to the wxWindowspeople, the wxPython mailing list, and so forth. You can also download wxPython there.
- Downloads, late developments, and documentation for the Python language can be found at the Python language Web site.
- *Programming Python* by Mark Lutz (O'Reilly, 1996), ISBN 1-56592-197-6, is the language reference I use for Python. It's more in a tutorial style rather than a reference, but so far it's the best I've found.
- If you want to make your brain hurt, you can read about Christian Tismer's Stackless Python, which efficiently implements coroutines, generators, continuations, and micro-threads. This is not for the faint of heart.
- The Vaults of Parnassus' Python resources are the best general starting place for information on text-based user interface tools in Python.
- Check out the JPython home page.
- Python.net is the development home for Python for .NET.
- Find out what happened at the Ninth International Python Conference in March.
- Read JPython: the Felicitous union of Python and Java, an excerpt from O'Reilly's book Learning Python.

This article was first published on IBM developerWorks and is re-printed with the permission of the author, author's agent and IBM developerWorks. The originals can be found at:

<http://www-106.ibm.com/developerworks/library/l-wxpy/?dwzone=linux?open&=amp;=amp;l=335,t=gr,p=wxPython>

Debian Conference – Australia,

February 4–5 2002

As part of the upcoming Linux.conf.au (see www.linux.org.au/conf) Australian National Linux Conference, a smaller Debian Conference is being planned for the preceding two days in Brisbane, Queensland, Australia. Linux.conf.au (LCA) is a 'roaming' conference under the auspices of Linux Australia and the regional Australian Linux User Groups. Last January 2001 it was held at the University of New South Wales in Sydney; prior to that was the Conference of Australian Linux Users in Melbourne in 1999.

Linux.conf.au is being held from Wednesday February 6th until Saturday February 9th 2002. Our aim is to hold a smaller sub-conference for Debian developers and users, focused specifically on Debian. However, so as not to detract from the LCA conference, we are hoping to take any prospective speakers from the Debian conference that may have a wider appeal and transplant them (and their presentation) to the main LCA conference.

Attendance to this is seen as an extension to the LCA conference, as LCA will be providing a venue for this event.

We are hoping to provide this event as an avenue for Debian Developers in the entire Asia-Pacific region (and beyond) to have an excuse to congregate in one place in the not too distant future. It is envisaged that the Debian Conference will run on Monday the 4th and Tuesday the 5th of February 2002. This announcement serves also as a Call For Participation: if you would like to present a topic of your choice at the Debian Conference, please email lca-cfp@linux.org.au (the main Call for Papers email address for LCA). Please note on your email that you would like to present at the Debian Conference. Please suggest several topics if you have more than one, with a brief summary of what you will be covering, etc.



The above information can be found at, and is updated at:



<http://www.linux.org.au/conf/debiancon.html>



For further information on LCA 2002, please email lca-organisers@linux.org.au. For further information on the Debian Conference, please email either jeb@debian.org, or ajt@debian.org.




AUUG 2001: The Conference Photos




Photographer: Andrew McRae (amcrae@cisco.com)











Shell Programming

Author: Katja and Guido Socher <katja@linuxfocus.org>,
<guido@linuxfocus.org>

ABSTRACT:

In this article we explain how to write little shell scripts and give many examples.

WHY SHELL PROGRAMMING?

Even though there are various graphical interfaces available for Linux the shell still is a very neat tool. The shell is not just a collection of commands but a really good programming language. You can automate a lot of tasks with it, the shell is very good for system administration tasks, you can very quickly try out if your ideas work which makes it very useful for simple prototyping and it is very useful for small utilities that perform some relatively simple tasks where efficiency is less important than ease of configuration, maintenance and portability. So let's see now how it works:

CREATING A SCRIPT

There are a lot of different shells available for Linux but usually the bash (bourne again shell) is used for shell programming as it is available for free and is easy to use. So all the scripts we will write in this article use the bash (but will most of the time also run with its older sister, the bourne shell). For writing our shell programs we use any kind of text editor, e.g. nedit, kedit, emacs, vi...as with other programming languages. The program must start with the following line (it must be the first line in the file):

```
#!/bin/sh
```

The #! characters tell the system that the first argument that follows on the line is the program to be used to execute this file. In this case /bin/sh is shell we use.

When you have written your script and saved it you have to make it executable to be able to use it. To make a script executable type

```
chmod +x filename
```

Then you can start your script by typing:

```
./filename
```

COMMENTS

Comments in shell programming start with # and go until the end of the line. We really recommend you to

use comments. If you have comments and you don't use a certain script for some time you will still know immediately what it is doing and how it works.

Variables

As in other programming languages you can't live without variables. In shell programming all variables have the datatype string and you do not need to declare them. To assign a value to a variable you write:

```
varname=value
```

To get the value back you just put a dollar sign in front of the variable:

```
#!/bin/sh
# assign a value:
a="hello world"
# now print the content of "a":
echo "A is:"
echo $a
```

Type this lines into your text editor and save it e.g. as first. Then make the script executable by typing chmod +x first in the shell and then start it by typing

```
./first
```

The script will just print:

```
A is:
hello world
```

Sometimes it is possible to confuse variable names with the rest of the text:

```
num=2
echo "this is the $numnd"
```

This will not print "this is the 2nd" but "this is the " because the shell searches for a variable called numnd which has no value. To tell the shell that we mean the variable num we have to use curly braces:

```
num=2
echo "this is the ${num}nd"
```

This prints what you want:

```
this is the 2nd
```

There are a number of variables that are always automatically set. We will discuss them further down when we use them the first time. If you need to handle mathematical expressions then you need to use programs such as expr (see table below). Besides the normal shell variables that are only valid within the shell program there are also environment variables. A variable preceded by the keyword export is an environment variable. We will not talk about them here any further since they are normally only used in login scripts.

SHELL COMMANDS AND CONTROL STRUCTURES

There are three categories of commands which can be used in shell scripts:

1) Unix commands:

Although a shell script can make use of any unix commands here are a number of commands which are more often used than others. These commands can generally be described as commands for file and text manipulation.

Command syntax	Purpose
echo "some text"	write some text on your screen
ls	list files
wc -l file wc -w file wc -c file	count lines in file or count words in file or count number of characters
cp sourcefile destfile	copy sourcefile to destfile
mv oldname newname	rename or move file
rm file	delete a file
grep 'pattern' file	search for strings in a file Example: grep 'searchstring' file.txt
cut -b colnum file	get data out of fixed width columns of text Example: get character positions 5 to 9 cut -b5-9 file.txt Do not confuse this command with "cat" which is something totally different
cat file.txt	write file.txt to stdout (your screen)
file somefile	describe what type of file somefile is
read var	prompt the user for input and write it into a variable (var)
sort file.txt	sort lines in file.txt
uniq	remove duplicate lines, used in combination with sort since uniq removes only duplicated consecutive lines Example: sort file.txt uniq
expr	do math in the shell Example: add 2 and 3 expr 2 "+" 3
find	search for files Example: search by name: find -name filename -print This command has many different possibilities and options. It is unfortunately too much to explain it all in this article.
tee	write data to stdout (your screen) and to a file Normally used like this: somecommand tee outfile It writes the output of somecommand to the screen and to the file outfile

basename file	return just the file name of a given name and strip the directory path Example: basename /bin/tux returns just tux
dirname file	return just the directory name of a given name and strip the actual file name Example: dirname /bin/tux returns just /bin
head file	print some lines from the beginning of a file
tail file	print some lines from the end of a file
sed	sed is basically a find and replace program. It reads text from standard input (e.g from a pipe) and writes the result to stdout (normally the screen). The search pattern is a regular expression (see references). This search pattern should not be confused with shell wildcard syntax. To replace the string linuxfocus with LinuxFocus in a text file use: cat text.file sed 's/linuxfocus/LinuxFocus/' > newtext file This replaces the first occurrence of the string linuxfocus in each line with LinuxFocus. If there are lines where linuxfocus appears several times and you want to replace all use: cat text.file sed 's/linuxfocus/LinuxFocus/g' > newtext file
awk	Most of the time awk is used to extract fields from a text line. The default field separator is space. To specify a different one use the option -F: cat file.txt awk -F, '{print \$1 " ", " \$3 }' Here we use the comma (,) as field separator and print the first and third (\$1 \$3) columns. If file.txt has lines like: Adam Bor, 34, India Kerry Miller, 22, USA then this will produce: Adam Bor, India Kerry Miller, USA There is much more you can do with awk but this is a very common use.

2) Concepts: Pipes, redirection and backtick

They are not really commands but they are very important concepts.

Pipes (|): send the output (stdout) of one program to the input (stdin) of another program.

```
grep "hello" file.txt | wc -l
```

finds the lines with the string hello in file.txt and then counts the lines. The output of the grep command is used as input for the wc command. You can concatenate as many commands as you like in that

way (within reasonable limits).

redirection: writes the output of a command to a file or appends data to a file

> writes output to a file and overwrites the old file in case it exists

>> appends data to a file (or creates a new one if it doesn't exist already but it never overwrites anything).

Backtick: The output of a command can be used as command line arguments (not stdin as above, command line arguments are any strings that you specify behind the command such as file names and options) for another command. You can as well use it to assign the output of a command to a variable.

The command:

```
find . -mtime -1 -type f -print
```

finds all files that have been modified within the last 24 hours (-mtime -2 would be 48 hours). If you want to pack all these files into a tar archive (file.tar) the syntax for tar would be:

```
tar xvf file.tar infile1 infile2 ...
```

Instead of typing it all in you can combine the two commands (find and tar) using backticks. Tar will then pack all the files that find has printed:

```
#!/bin/sh
# The ticks are backticks (') not normal quotes
('):
tar -zcvf lastmod.tar.gz `find . -mtime -1 -type f
-print`
```

3) Control structures

The "if" statement tests if the condition is true (exit status is 0, success). If it is the "then" part gets executed:

```
if ....; then
    ....
elif ....; then
    ....
else
    ....
fi
```

Most of the time a very special command called test is used inside if-statements. It can be used to compare strings or test if a file exists, is readable etc...

The "test" command is written as square brackets "[]". Note that space is significant here: Make sure that you always have space around the brackets. Examples:

```
[ -f "somefile" ] : Test if somefile is a file.
[ -x "/bin/ls" ]  : Test if /bin/ls exists and is
executable.
[ -n "$var" ]     : Test if the variable $var
contains something
[ "$a" = "$b" ]   : Test if the variables "$a"
and "$b" are equal
```

Run the command "man test" and you get a long list

of all kinds of test operators for comparisons and files. Using this in a shell script is straight forward:

```
#!/bin/sh
if [ "$SHELL" = "/bin/bash" ]; then
    echo "your login shell is the bash (bourne again
shell)"
else
    echo "your login shell is not bash but $SHELL"
fi
```

The variable \$SHELL contains the name of the login shell and this is what we are testing here by comparing it against the string

```
"/bin/bash"
```

Shortcut operators: People familiar with C will welcome the following expression:

```
[ -f "/etc/shadow" ] && echo "This computer uses
shadow passwords"
```

The && can be used as a short if-statement. The right side gets executed if the left is true. You can read this as AND. Thus the example is: "The file /etc/shadow exists AND the command echo is executed". The OR operator (||) is available as well.

Here is an example:

```
#!/bin/sh
mailfolder=/var/spool/mail/james
[ -r "$mailfolder" ] || { echo "Can not read
$mailfolder" ; exit 1; }
echo "$mailfolder has mail from:"
grep "From" $mailfolder
```

The script tests first if it can read a given mailfolder. If yes then it prints the "From" lines in the folder. If it cannot read the file \$mailfolder then the OR operator takes effect. In plain English you read this code as "Mailfolder readable or exit program". The problem here is that you must have exactly one command behind the OR but we need two:

- print an error message
- exit the program

To handle them as one command we can group them together in an anonymous function using curly braces. Functions in general are explained further down. You can do everything without the ANDs and ORs using just if-statements but sometimes the shortcuts AND and OR are just more convenient. The case statement can be used to match (using shell wildcards such as * and ?) a given string against a number of possibilities.

```
case ... in
...) do something here;;
esac
```

Let's look at an example. The command file can test what kind of filetype a given file is:

```
file lf.gz
```

returns:

```
lf.gz: gzip compressed data, deflated, original
```

```
filename,  
last modified: Mon Aug 27 23:09:18 2001, os: Unix
```

We use this now to write a script called [27]smartzip that can uncompress bzip2, gzip and zip compressed files automatically :

```
#!/bin/sh  
ftype=`file "$1"`  
case "$ftype" in  
"$1: Zip archive"*)  
    unzip "$1" ;;  
"$1: gzip compressed"*)  
    gunzip "$1" ;;  
"$1: bzip2 compressed"*)  
    bunzip2 "$1" ;;  
*) error "File $1 can not be uncompressed with  
smartzip";;  
esac
```

Here you notice that we use a new special variable called \$1. This variable contains the first argument given to a program. Say we run:

```
smartzip articles.zip
```

then \$1 will contain the string *articles.zip*

The select statement is a bash specific extension and is very good for interactive use. The user can select a choice from a list of different values:

```
select var in ... ; do  
    break  
done
```

.... now \$var can be used

Here is an example:

```
#!/bin/sh  
echo "What is your favourite OS?"  
select var in "Linux" "Gnu Hurd" "Free  
BSD" "Other"; do  
    break  
done  
echo "You have selected $var"
```

Here is what the script does:

```
What is your favourite OS?  
1) Linux  
2) Gnu Hurd  
3) Free BSD  
4) Other  
#? 1  
You have selected Linux
```

In the shell you have the following loop statements available:

```
while ...; do  
    ....  
done
```

The while-loop will run while the expression that we test for is true. The keyword "break" can be used to leave the loop at any point in time. With the keyword "continue" the loop continues with the next iteration and skips the rest of the loop body. The for-loop takes a list of strings (strings separated by space) and assigns them to a variable:

```
for var in ...; do  
    ....  
done
```

The following will e.g. print the letters A to C on the screen:

```
#!/bin/sh  
for var in A B C ; do  
    echo "var is $var"  
done
```

A more useful example script, called showrpm, prints a summary of the content of a number of RPM-packages:

```
#!/bin/sh  
# list a content summary of a number of RPM  
packages  
# USAGE: showrpm rpmfile1 rpmfile2 ...  
# EXAMPLE: showrpm /cdrom/RedHat/RPMS/*.rpm  
for rpmpackage in $*; do  
    if [ -r "$rpmpackage" ];then  
        echo "===== $rpmpackage  
===== "  
        rpm -qi -p $rpmpackage  
    else  
        echo "ERROR: cannot read file $rpmpackage"  
    fi  
done
```

Above you can see the next special variable, \$* which contains all the command line arguments. If you run:

```
showrpm openssh.rpm w3m.rpm webgrep.rpm
```

then \$* contains the 3 strings openssh.rpm, w3m.rpm and webgrep.rpm. The GNU bash knows until-loops as well but generally while and for loops are sufficient.

QUOTING

Before passing any arguments to a program the shell tries to expand wildcards and variables. To expand means that the wildcard (e.g. *) is replaced by the appropriate file names or that a variable is replaced by its value. To change this behaviour you can use quotes: Let's say we have a number of files in the current directory. Two of them are jpg-files, mail.jpg and tux.jpg.

```
#!/bin/sh  
echo *.jpg
```

This will print "mail.jpg tux.jpg".

Quotes (single and double) will prevent this wildcard expansion:

```
#!/bin/sh  
echo "*.jpg"  
echo '*.jpg'
```

This will print "*.jpg" twice. Single quotes are most strict. They prevent even variable expansion. Double quotes prevent wildcard expansion but allow variable expansion:

```
#!/bin/sh
```

```
echo $SHELL
echo "$SHELL"
echo '$SHELL'
```

This will print:

```
/bin/bash
/bin/bash
$SHELL
```

Finally there is the possibility to take the special meaning of any single character away by preceding it with a backslash:

```
echo \*.jpg
echo \$SHELL
```

This will print:

```
*.jpg
$SHELL
```

HERE DOCUMENTS

Here documents are a nice way to send several lines of text to a command. It is quite useful to write a help text in a script without having to put echo in front of each line. A "Here document" starts with << followed by some string that must also appear at the end of the here document. Here is an [28]example script, called ren, that renames multiple files and uses a here document for its help text:

```
#!/bin/sh
# we have less than 3 arguments. Print the help
text:
if [ $# -lt 3 ] ; then
cat <<HELP
ren -- renames a number of files using sed regular
expressions

USAGE: ren 'regexp' 'replacement' files...

EXAMPLE: rename all *.HTM files in *.html:
ren 'HTM$' 'html' *.HTM

HELP
exit 0
fi
OLD="$1"
NEW="$2"
# The shift command removes one argument from the
list of
# command line arguments.
shift
shift
# $* contains now all the files:
for file in $*; do
if [ -f "$file" ] ; then
newfile=`echo "$file" | sed
"s/${OLD}/${NEW}/g`"
if [ -f "$newfile" ] ; then
echo "ERROR: $newfile exists already"
else
echo "renaming $file to $newfile ..."
mv "$file" "$newfile"
fi
fi
done
```

This is the most complex script so far. Let's discuss it a little bit. The first if-statement tests if we have provided at least 3 command line parameters. (The special variable \$# contains the number of arguments.) If not, the help text is sent to the command cat which in turn sends it to the screen. After printing the help text we exit the program. If there are 3 or more arguments we assign the first argument to the variable OLD and the second to the

variable NEW. Next we shift the command line parameters twice to get the third argument into the first position of \$*. With \$* we enter the for loop. Each of the arguments in \$* is now assigned one by one to the variable \$file. Here we first test that the file really exists and then we construct the new file name by using find and replace with sed. The backticks are used to assign the result to the variable newfile. Now we have all we need: The old file name and the new one. This is then used with the command mv to rename the files.

FUNCTIONS

As soon as you have a more complex program you will find that you use the same code in several places and also find it helpful to give it some structure. A function looks like this:

```
functionname()
{
# inside the body $1 is the first argument given
to the function
# $2 the second ...
body
}
```

You need to "declare" functions at the beginning of the script before you use them. Here is a script called xtitlebar which you can use to change the name of a terminal window.

```
#!/bin/sh
# vim: set sw=4 ts=4 et: help()
{
cat <<HELP
xtitlebar -- change the name of an xterm, gnome-
terminal or kde konsole USAGE: xtitlebar [-h]
"string_for_titlebar" OPTIONS: -h this help text
EXAMPLE: xtitlebar "cvs" HELP
exit 0
} [ -z "$1" ] && help
[ "$1" = "-h" ] && help # sent the escape sequence
to change the xterm title bar:
echo -e "\033]0;$1\007"
#
```

If you have several of them open it is easier to find them. The script sends an escape sequence which is interpreted by the terminal and causes it to change the name in the titlebar. The script uses a function called help. As you can see the function is defined once and then used twice:

```
#!/bin/sh
# vim: set sw=4 ts=4 et:

help()
{
cat <<HELP
xtitlebar -- change the name of an xterm, gnome-
terminal or kde konsole

USAGE: xtitlebar [-h] "string_for_titelbar"

OPTIONS: -h help text

EXAMPLE: xtitlebar "cvs"

HELP
exit 0
}

# in case of error or if -h is given we call the
function help:
[ -z "$1" ] && help
[ "$1" = "-h" ] && help

# send the escape sequence to change the xterm
```

```

titelbar:
echo -e "\033]0;${1}\007"
#

```

It's a good habit to always have extensive help inside the scripts. This makes it possible for others (and you) to use and understand the script.

COMMAND LINE ARGUMENTS

We have seen that \$* and \$1, \$2 ... \$9 contain the arguments that the user specified on the command line (The strings written behind the program name). So far we had only very few or rather simple command line syntax (a couple of mandatory arguments and the option -h for help). But soon you will discover that you need some kind of parser for more complex programs where you define your own options. The convention is that all optional parameters are preceded by a minus sign and must come before any other arguments (such as e.g file names). There are many possibilities to implement a parser. The following while loop combined with a case statement is a very good solution for a [30]generic parser:

```

#!/bin/sh
help()
{
    cat <<HELP
    Write usage and help text here
    HELP
    exit 0
}
while [ -n "$1" ]; do
case $1 in
    -h) help;shift 1;; # function help is called
    -f) opt_f=1;shift 1;; # variable opt_f is set
    -l) opt_l=$2;shift 2;; # -l takes an argument
-> shift by 2
    --) shift;break;; # end of options
    *) echo "error: no such option $1. -h for
help";exit 1;;
    *) break;;
esac
done
echo "opt_f is $opt_f"
echo "opt_l is $opt_l"
echo "first arg is $1"
echo "2nd arg is $2"

```

Try it out! You can run it e.g with:

```
cmdparser -l hello -f -- -somefile1 somefile2
```

It produces:

```

opt_f is 1
opt_l is hello
first arg is -somefile1
2nd arg is somefile2

```

How does it work? Basically it loops through all arguments and matches them against the case statement. If it finds a matching one it sets a variable and shifts the command line by one. The unix convention is that options (things starting with a minus) must come first. You may indicate that this is the end of option by writing two minus signs (--). You need it e.g with grep to search for a string starting with a minus sign:

```

Search for -xx- in file f.txt:
grep -- -xx- f.txt

```

Our option parser can handle the -- too as you can see in the listing above.

EXAMPLES

A general purpose selecton

Now we have discussed almost all components that you need to write a script. All good scripts should have help and you can as well have our generic option parser even if the script has just one option. Therefore it is a good idea to have a dummy script, called framework.sh,

```

#!/bin/sh
# vim: set sw=4 ts=4 et:
help()
{
    cat <<HELP
    Write usage and help text here
    HELP
    exit 0
}
error()
{
    # print an error and exit
    echo "$1"
    exit 1
}
# The option parser, change it as needed
# In this example -f and -h take no arguments -l
# takes an argument
# after the l
while [ -n "$1" ]; do
case $1 in
    -h) help;shift 1;; # function help is called
    -f) opt_f=1;shift 1;; # variable opt_f is set
    -l) opt_l=$2;shift 2;; # -l takes an argument
-> shift by 2
    --) shift;break;; # end of options
    *) echo "error: no such option $1. -h for
help";exit 1;;
    *) break;;
esac
done
# The main program of you script comes after
this line
error "ERROR: This is a generic script framework
you must modify it first"
#

```

which you can use as a framework for other scripts. If you want to write a new script you just make a copy:

```
cp framework.sh myscript
```

and then insert the actual functionality into "myscript".

Let's now look at two more examples:

A binary to decimal number converter

The script b2d:

```

#!/bin/sh
# vim: set sw=4 ts=4 et:
help()
{
    cat <<HELP
    b2h -- convert binary to decimal USAGE: b2h [-h]
    binarynum OPTIONS: -h this help EXAMPLE: b2h
    111010
    will return 58
    HELP
    exit 0
}
error()
{
    # print an error and exit

```

```

    echo "$1"
    exit 1
} lastchar()
# return the last character of a string in
$rvall
if [ -z "$1" ]; then
    # empty string
    rval=""
    return
fi
# wc puts some space behind the output this is
why we need sed:
numofchar='echo -n "$1" | wc -c | sed 's/ //g'
# now cut out the last char
rval='echo -n "$1" | cut -b $numofchar'
} chop()
# remove the last character in string and
return it in $rval
if [ -z "$1" ]; then
    # empty string
    rval=""
    return
fi
# wc puts some space behind the output this is
why we need sed:
numofchar='echo -n "$1" | wc -c | sed 's/ //g'
if [ "$numofchar" = "1" ]; then
    # only one char in string
    rval=""
    return
fi
numofcharminus1='expr $numofchar "-" 1'
# now cut all but the last char:
rval='echo -n "$1" | cut -b
0-${numofcharminus1}'
}
while [ -n "$1" ]; do
case $1 in
-h) help;shift 1;; # function help is called
--) shift;break;; # end of options
-*) error "error: no such option $1. -h for
help";;
*) break;;
esac
done # The main program
sum=0
weight=1
# one arg must be given:
[ -z "$1" ] && help
binnum="$1"
binnumorig="$1" while [ -n "$binnum" ]; do
    lastchar "$binnum"
    if [ "$rval" = "1" ]; then
        sum='expr "$weight" "+" "$sum"'
    fi
    # remove the last position in $binnum
    chop "$binnum"
    binnum="$rval"
    weight='expr "$weight" "*" 2'
done echo "binary $binnumorig is decimal $sum"
#

```

converts a binary number (e.g 1101) into its decimal equivalent. It is an example that shows that you can do simple mathematics with expr:

```

#!/bin/sh
# vim: set sw=4 ts=4 et:
help()
{
    cat <<HELP
b2h -- convert binary to decimal
USAGE: b2h [-h] binarynum
OPTIONS: -h help text
EXAMPLE: b2h 111010
will return 58
HELP
    exit 0
}
error()
{
    # print an error and exit
    echo "$1"
    exit 1
}

```

```

lastchar()
{
    # return the last character of a string in
$rvall
if [ -z "$1" ]; then
    # empty string
    rval=""
    return
fi
# wc puts some space behind the output this is
why we need sed:
numofchar='echo -n "$1" | wc -c | sed 's/ //g'
# now cut out the last char
rval='echo -n "$1" | cut -b $numofchar'
}
chop()
# remove the last character in string and
return it in $rval
if [ -z "$1" ]; then
    # empty string
    rval=""
    return
fi
# wc puts some space behind the output this is
why we need sed:
numofchar='echo -n "$1" | wc -c | sed 's/ //g'
if [ "$numofchar" = "1" ]; then
    # only one char in string
    rval=""
    return
fi
numofcharminus1='expr $numofchar "-" 1'
# now cut all but the last char:
rval='echo -n "$1" | cut -b
0-${numofcharminus1}'
}
while [ -n "$1" ]; do
case $1 in
-h) help;shift 1;; # function help is called
--) shift;break;; # end of options
-*) error "error: no such option $1. -h for
help";;
*) break;;
esac
done
# The main program
sum=0
weight=1
# one arg must be given:
[ -z "$1" ] && help
binnum="$1"
binnumorig="$1"
while [ -n "$binnum" ]; do
    lastchar "$binnum"
    if [ "$rval" = "1" ]; then
        sum='expr "$weight" "+" "$sum"'
    fi
    # remove the last position in $binnum
    chop "$binnum"
    binnum="$rval"
    weight='expr "$weight" "*" 2'
done
echo "binary $binnumorig is decimal $sum"
#

```

The algorithm used in this script takes the decimal weight (1,2,4,8,16...) of each digit starting from the right most digit and adds it to the sum if the digit is a 1. Thus "10" is: $0 * 1 + 1 * 2 = 2$

To get the digits from the string we use the function lastchar. This uses wc -c to count the number of characters in the string and then cut to cut out the last character. The chop function has the same logic but removes the last character, that is it cuts out everything from the beginning to the character before the last one.

A FILE ROTATION PROGRAM

Perhaps you are one of those who save all outgoing mail to a file. After a couple of months this file becomes rather big and it makes the access slow if you load it into your mail program. The following script **rotatefile** can help you. It renames the mailfolder, let's call it outmail, to outmail.1 if there was already an outmail.1 then it becomes outmail.2 etc...

```
#!/bin/sh
# vim: set sw=4 ts=4 et:
ver="0.1"
help()
{
    cat <<HELP
rotatefile -- rotate the file name
USAGE: rotatefile [-h] filename
OPTIONS: -h help text
EXAMPLE: rotatefile out
This will e.g rename out.2 to out.3, out.1 to
out.2, out to out.1
and create an empty out-file
The max number is 10
version $ver
HELP
    exit 0
}
error()
{
    echo "$1"
    exit 1
}
while [ -n "$1" ]; do
case $1 in
    -h) help; shift 1;;
    --) break;;
    *) echo "error: no such option $1. -h for
help"; exit 1;;
    *) break;;
esac
done
# input check:
if [ -z "$1" ]; then
    error "ERROR: you must specify a file, use -h for
help"
fi
filen="$1"
# rename any .1 , .2 etc file:
for n in 9 8 7 6 5 4 3 2 1; do
    if [ -f "$filen.$n" ]; then
        p='expr $n + 1'
        echo "mv $filen.$n $filen.$p"
        mv $filen.$n $filen.$p
    fi
done
# rename the original file:
if [ -f "$filen" ]; then
    echo "mv $filen $filen.1"
    mv $filen $filen.1
fi
echo touch $filen
touch $filen
```

How does the program work? After checking that the user provided a filename we go into a for loop counting from 9 to 1. File 9 is now renamed to 10, file 8 to 9 and so on. After the loop we rename the original file to 1 and create an empty file with the name of the original file.

DEBUGGING

The most simple debugging help is of course the command echo. You can use it to print specific variables around the place where you suspect the mistake. This is probably what most shell programmers use 80% of the time to track down a mistake. The advantage of a shell script is that it does not require any re-compilation and inserting an "echo" statement is done very quickly. The shell has a real debug mode as well. If there is a mistake in your script "strangescript" then you can debug it like this:

```
sh -x strangescript
```


This will execute the script and show all the statements that get executed with the variables and wildcards already expanded. The shell also has a mode to check for syntax errors without actually executing the program. To use this run:

```
sh -n your_script
```

If this returns nothing then your program is free of syntax errors. We hope you will now start writing your own shell scripts. Have fun!

REFERENCES

- The man page and the info page of bash are very good and you will discover many more tricks there: man bash tkinfo bash (I hate the standard info browser but [34]tkinfo or konqueror (the url is info:/bash/Top) are quite good)
- The LinuxFocus article on GNU file utilities (<http://mercury.chem.pitt.edu/~tiho/LinuxFocus/English/November2000/article131.shtml>)
- The LinuxFocus article on Regular expressions (<http://mercury.chem.pitt.edu/~tiho/LinuxFocus/English/July1998/article53.html>)
- The LinuxFocus article on AWK (<http://mercury.chem.pitt.edu/~tiho/LinuxFocus/English/September1999/article103.html>)
- The LinuxFocus article on Basic UNIX commands (<http://mercury.chem.pitt.edu/~tiho/LinuxFocus/English/November1998/article68.html>)
- If bash isn't installed on your system or you want to get the newest version you can download it from the GNU website at <http://www.gnu.org> (to find out if bash is already installed on your system just type bash. If you don't get an error message bash is already installed)



**10% DISCOUNT
TO AUUG MEMBERS
ON OUR COMPLETE RANGE
OF COMPUTER, BUSINESS
AND GENERAL BOOKS**

**AMERICAN
BOOK STORE**

173 Elizabeth St, Brisbane Queensland 4000

Ph: (07) 3229 4677 Fax: (07) 3221 2171 Qld Country Freecall: 1800 177 395

american_bookstore@compuserve.com

Name: _____ Date: _____

Address: _____

Post Code: _____

Phone Number: _____

Payment Method: Cheque Money Order Amex Bankcard
 Diners Mastercard Visa

Card Number: _____

Expiry Date: _____ Signature: _____

This is a: Special Order Mail Order Book on Hold

QUANTITY	TITLE	PRICE
.....
.....
.....
.....
.....
.....
.....
.....
.....
.....

	SUBTOTAL	\$
	LESS 10% DISCOUNT	\$
	POST & PACK	\$
	TOTAL	\$

POSTAGE AND HANDLING FEES: 1 BOOK \$6.00 2-4 BOOKS \$7.00
 BOOKS OVER \$70.00 WE WILL SEND CERTIFIED - PLEASE ADD ANOTHER \$1.50 OR WAIVE CERTIFIED DELIVERY.

FOR SPECIAL ORDERS, PLEASE ENCLOSE \$10.00 PER BOOK AS A DEPOSIT.

AUUG Security Symposium 2001

19-21 November 2001 - Brisbane
Auditorium of the Primary Industries,
Department, Ann St, Brisbane

INTRODUCTION

The AUUG Security Symposium provides a forum for discussion of security technologies, techniques and management. Our society today is highly dependant on our almost pervasively interconnected systems. Hence we are also dependant upon the security of these systems. As Governments and private industry become increasingly aware of the vulnerability of our systems there is a growing requirement for security education and for practioners to share their knowledge for the greater good.

This symposium aims to fill a gap in the Australian conference scene between the high cost commercial conferences where attendees hear mainly marketing pitches and the academic-based research conferences. It is unashamedly for the practioner in the field who wants to share (or know about) how to secure their systems (be it a PC operating systems, a huge network or a client server application).

The symposium will be a three day event with paper presentations and tutorials. This will be the best opportunity in Australia this year to meet, discuss and debate your ideas and experiences on information security.

ACCOMMODATION

Carlton Crest Hotel
King George Square
(cnr Ann & Roma Street)
Brisbane QLD 4000

is offering the following accommodation rates for delegates attending the symposium. Please state that you are attending the "AUUG Security Symposium" when booking to obtain these rates.

Carlton Tower - Carlton Deluxe \$148.00 per room,
per night

Crest Tower - Crest Standard \$125.00 per room,
per night

Please note that all prices are inclusive of 10% GST.
Telephone : + 61 7 3229 9111
Toll Free : 1800 777 123 (Aust. Only)
Facsimile : + 61 7 3229 9618

These rates are for single, double or twin style accommodation and are inclusive of the rooftop gymnasium, sauna and pool. Full buffet breakfast is available in the Spice Market Restaurant 7 days a week. The breakfast rate is \$18.00 per person per

day. Delegates are responsible for settling their own accounts on departure, directly with the Carlton Crest.

Please note: Any amendments or cancellations to hotel bookings must be made directly with the Carlton Crest. Limited car parking is available beneath the Carlton Crest Hotel. For guests attending functions or staying in-house, parking is available at a flat rate of \$12.00 per vehicle per day. Parking is also conveniently located at the King George Square Car park adjacent to the hotel.

SECURITY SYMPOSIUM CONTACT INFORMATION

The Security Symposium committee is listed below:

Committee

Gary Gaskell (gary.gaskell@member.sage-au.org.au)
Level 8, 229 Elizabeth St, Brisbane
IT Department
Bank of Queensland
Phn: +61 7 3212 3447
Mob: 0419 714 191
Fax: +61 7 3212 3229

Warren Toomey (wkt@tuhs.org)
School of IT
Bond University
Work: +61 7 5595 3331

Duncan Unwin (Duncan.Unwin@qsipayments.com)
QSI Payments, Inc.
22/300 Adelaide St
BRISBANE, QLD, AUSTRALIA 4000
Phn: +61-7-3224 9856
Mob: +61-419-020 485
Fax: +61-7-3210 2566

Lawrie Brown (Lawrie.Brown@canb.auug.org.au)
School of Computer Science
Australian Defence Force Academy

INFORMATION DISCLAIMER

The speakers, topics and times are correct at the time of publishing. In the event of unforeseen circumstances, AUUG reserves the right to alter or delete items from the Symposium Programme. Each Delegate named herein, or their later substitute, accepts that AUUG, their agents, speakers, chairmen, secretariat or any other persons ("the parties") involved in the preparation of this pamphlet or in the planning or presentation of the Symposium, do not accept any contractual, tortious or other form of liability for loss or damage suffered by the delegate or their later substitute relying on any statement representation advice or opinion (whether true or otherwise, written or oral) and whether due to the negligence of any of the said parties by this disclaimer of liability to exclude liability, if any, for such statement representation advice or opinion, and that the said party may rely on this disclaimer of liability

in the event of any demand claim suit or action brought against any or all of them arising out of any statement representation advice or opinion. The authors have prepared this material for Professional Development purposes. Although they trust that it will be useful for this purpose, neither the authors nor AUUG can warrant that the use of this material would be adequate to discharge the legal or professional liability of members in the conduct of their practices. AUUG reserves the right to cancel the event.

For enquiries on conference registration, accommodation arrangements, promotion, venue and other matters not relating to the submission of papers, contact the AUUG Business Manager busmgr@auug.org.au or phone 02 8824 9511 or 1800 625 655 or Fax (02) 8824 9522.

PRELIMINARY PROGRAM

Monday 19 November	Tutorials	See the tutorials page
Tuesday 20 November		
8:00 – 9:00	Registration & coffee	
9:00 – 9:10	Welcome	
9:10 – 10:00	Keynote	Adrian McCullagh, Spyrus Inc.
	Coffee Break	
10:30 – 11:00	New trends in authenticating payments on-line	Duncan Unwin, QSI Payments
11:00 – 12:00	Advanced log analysis techniques	M. I. Cohen, A. Corby & T. Kaiser, DSD
	Lunch (provided)	
13:30 – 14:00	Information security, the Australian Privacy regime, and what it means for IT security practitioners	Brian Denehy and Bernard Hill, 90East (Asia Pacific) Pty Ltd
14:00 – 14:30	Wireless security	
14:30 – 15:00	Management issues in IDS	Nathan Carey, Information Security Research Centre, QUT
	Coffee break	
15:45 – 16:15	Financial Sector Security Issues	Steven Anderssen, Commonwealth Bank
16:15 – 16:45	Wireless Insecurity	Neal Wise, esec

16:45 – 17:00	Wrap up & notices	
7:00	eGlobal eSecurity Symposium Dinner	
Wednesday 21 November		
8:30	Coffee	
9:00 – 9:10	Plenary/Notices	
9:10 – 10:00	Guest speaker	Russ Cooper, NTbugtraq editor and TruSecure/eGlobal
	Coffee break	
10:45 – 11:15	Securing NFS in a Teaching Laboratory environment	Bob Edwards and Matt Pratt, ANU
11:15 – 11:45	Holistic Approach to Information Security – Management & Assurance	Mark Ames, eGlobal eSecurity Australia
11:45 – 12:15	Sun Security Features	Eric Halil, Sun Microsystems
	Lunch (provided)	
14:00 – 14:30	Network Attacks: Trends in Australia and Overseas	Jamie Gillespie, AusCERT
14:30 – 15:00	Australian Security and Legal Issues	Darrell Betts, AFP
15:30 – 16:45	Panel: Ask the experts!	Danny Smith, Lawrie Brown, Mark McPherson, Darrell Betts
16:45 – 17:00	Wrap up	

PRELIMINARY TUTORIAL PROGRAMME

Monday 19 November	Time	Tut No.	Tutorial Title and Outline
Half Day Morning Tutorial	9am – 12.30p m	M1	<p>Firewalls, by Lawrie Brown</p> <p>With the ever increasing growth and pervasiveness of the Internet, more and more organisations find that they need to connect to the Internet in order to fulfil their goals. However, there are persistent security concerns with such a connection. The usual approach to reducing these concerns is to install a firewall to provide perimeter defence around private networks which supplies a single controlled and monitored point of connection. The design, installation, and ongoing management of a firewall though, is a non-trivial task. This workshop will provide an overview of this process. This starts with the determination of an appropriate security policy, and then the specification of services to be supported and policy applied. From this a suitable firewall architecture can be selected from the range available, specific equipment chosen and configured. Then there is the ongoing management of the firewall, maintaining its safe configuration, responding to security events, and monitoring its ongoing use. The workshop will not discuss particular products, rather it is aimed to assist those who need to manage this process.</p> <p>Contents</p> <ul style="list-style-type: none"> • Gateway Policy • Gateway Design • Installation and Configuration • Gateway Management • Summary • More Resources • Appendices
Half Day Morning Tutorial	9am – 12.30p m	M2	<p>Practical IPSEC, by Adrian Close</p> <p>Networks on the Internet are increasingly turning to firewalls as a means of protecting themselves against external network-based attacks, creating their own small islands of trust.</p> <p>However, the increasing need for secure, inter-network communications requires extending that trust across the Internet itself – a risky proposition in an increasingly hostile network environment.</p> <p>Implementing IPSEC is one plausible solution and this tutorial will cover the fundamentals of doing this in the real world.</p> <p>Practical demonstrations of the technology involved will be given throughout the tutorial, which will include debugging techniques useful for successful deployment and interoperability of various IPSEC implementations.</p> <p>Topics:</p> <ul style="list-style-type: none"> • Why IPSEC? • Basic IPSEC – ESP, AH, SAs and SPIs. • Encryption algorithms – choices and availability • The problem of key exchange. • ISAKMP overview. • ISAKMP authentication using shared secrets and certificates. • PKI – myths and realities. • Alternatives to ISAKMP. • IPSEC implementations and interoperability issues. • IPSEC and IPV6 – a vision of the future.

Half Day Morning Tutorial	9am – 12.30p m	M3	<p>Securing Solaris, by John Warburton</p> <p>Solaris is currently the most widely implemented proprietary UNIX on the Internet today. Like all systems, the standard installation can have security issues – from out of date applications to insecure defaults.</p> <p>This tutorial will give the student an overview of the tasks required to secure a Solaris system in a number of environments – from workstation and server to firewall and web server.</p> <p>Main topics discussed will be:</p> <ul style="list-style-type: none"> • Base Operating System Installation • Operating System Hardening • Securing the Network Connection • Intrusion Detection • System Specifics: <ul style="list-style-type: none"> • Workstation • Server • Firewall • Web Server
Half Day Afternoon Tutorial	1:30p m – 5pm	M4	<p>A Cryptography Primer, by Lawrie Brown</p> <p>Data encryption algorithms form an important technical component in providing secure and authenticated electronic security and communications. This workshop is designed to provide attendees with a brief overview of the field of cryptography, the terms, techniques, and algorithms. It starts by introducing the classical cryptographic techniques which form the foundations of the field. We then survey modern private key ciphers, widely used for bulk and link data encryption, from DES to the new AES encryption algorithm Rijndael. Next we consider public key encryption algorithms and signature schemes, essential for the use of cryptography in large scale, wide area communications. We conclude with a brief look at a couple of cryptographic applications, illustrating the different ways these components are combined to build a security solution.</p> <p>Contents</p> <ul style="list-style-type: none"> • Introduction • Classical Cryptographic Techniques • Modern Cryptographic Algorithms • Private-Key Encryption Algorithms • Block Ciphers • Stream Ciphers and the Vemam cipher • Public-Key Encryption Algorithms • Digital Signature Algorithms • Hashing Functions • Key Management • Security in Practice – Secure Email • Security in Practice – SNMP • Security in Practice – One Time Pass words • Review • Terminology

Half Day Afternoon	1:30p m – 5pm	M5	<p>Building an open source firewall, by Michael Paddon</p> <p>Most network administrators today recognise the need for firewalls to enforce policy on inter-network traffic. While there are many commercial products addressing this need, many organisations are turning to open source solutions for reasons of performance, strength, robustness, transparency and price.</p> <p>One popular and mature open source firewall package is Darren Reed's ipfilter, which is available for a wide range of unix like systems, both open source and commercial. Attendees will learn, in detail, how to configure ipfilter for a wide range of real world situations and protocols. Topics covered include:</p> <ul style="list-style-type: none"> • introduction to the ipfilter processing model • filtering raw IP packets • filtering UDP and TCP packets • stateful packet filtering • effective ruleset structures • packet logging • customising icmp returns • tuning for performance • rule groups • complex protocols • authorisation • network address translation • general firewall configuration tips <p>Attendees are expected to be familiar with the IP, UDP and TCP protocols and should have some systems/network administration experience.</p>
Half Day Afternoon	1:30p m – 5pm	M6	<p>Secure Router Configuration, by Philip Smith, Cisco Systems</p>

Half Day Afternoon	1:30p m – 5pm	M7	<p>Securing BIND, by Pauline van Winsen</p> <p>The Domain Name System, DNS, is used to resolve names to IP addresses and vice versa and thus is one of the fundamental infrastructure protocols in use on the Internet and within organisations. Unfortunately, a number of security exploits have been revealed in specific versions of BIND and many default installations of BIND are exploitable.</p> <p>This tutorial will give the student an overview of the tasks required to secure a BIND name server and steps which can be taken to minimize the impact of any future vulnerabilities which may be found in BIND. Example configurations will be given for versions 8 and 9 of BIND.</p> <p>The main topics discussed will be:</p> <ul style="list-style-type: none"> • Obtaining and installing the latest BIND software • Improving BIND security through configuration options • Implementing BIND in a chroot environment • Running multiple versions of BIND on a single server to implement split-DNS • Future BIND security directions with zone and request authentication • Extending principles used to secure BIND to other crucial services • Alternative DNS implementations
-----------------------	---------------------	----	---

Chapter News Items Needed!

The AUUGN Editor Needs YOU!

Please submit any notices or information pertaining to our local chapter here

If you would like your local AUUG Chapter news and reviews to be listed here, send mail to auugn@auug.org.au

AUUG Chapter Meetings and Contact Details

CITY	LOCATION	OTHER
BRISBANE	Inn on the Park 507 Coronation Drive Toowong	For further information, contact the QAUUG Executive Committee via email (qauug-exec@auug.org.au). The techno-logically deprived can contact Rick Stevenson on (07) 5578-8933. To subscribe to the QAUUG announcements mailing list, please send an e-mail message to: <majordomo@auug.org.au> containing the message "subscribe qauug <e-mail address>" in the e-mail body.
CANBERRA	Australian National University	
HOBART	University of Tasmania	
MELBOURNE	Various. For updated information See: http://www.vic.auug.org.au/auugvic/av_meetings.html	The meetings alternate between Technical presentations in the odd numbered months and purely social occasions in the even numbered months. Some attempt is made to fit other AUUG activities into the schedule with minimum disruption.
PERTH	The Victoria League 276 Onslow Road Shenton Park	Meeting commences at 6.15pm
SYDNEY	TBA	

FOR UP-TO-DATE DETAILS ON CHAPTERS AND MEETINGS, INCLUDING THOSE IN ALL OTHER AUSTRALIAN CITIES, PLEASE CHECK THE AUUG WEBSITE AT [HTTP://WWW.AUUG.ORG.AU](http://www.auug.org.au) OR CALL THE AUUG OFFICE ON 1-800-625655.

Application for Institutional Membership

Section A: MEMBER DETAILS

The primary contact holds the full member voting rights and two designated representatives will be given membership rates to AUUG activities including chapter activities. In addition to the primary and two representatives, additional representatives can be included at a rate of \$88 each. **Please attach a separate sheet with details of all representatives to be included with your membership.**

NAME OF ORGANISATION: _____

Primary Contact

Surname _____ First Name _____
 Title: _____ Position _____
 Address _____
 Suburb _____ State _____ Postcode _____
 Telephone: Business _____ Facsimile _____
 Email _____ Local Chapter Preference _____

Section B: MEMBERSHIP INFORMATION.

Renewal/New Institutional Membership of AUUG \$429.00
(including Primary and Two Representatives)
 Surcharge for International Air Mail \$132.00
 Additional Representatives Number @ \$88.00

Rates valid as at 1 March 2000. Memberships valid through to 30 June 2001 and include 10% GST.

Section C: PAYMENT

Cheques to be made payable to AUUG Inc (Payment in Australian Dollars only)

For all overseas applications, a bank draft drawn on an Australian bank is required.
Please do not send purchase orders.

-OR-

Please debit my credit card for AS\$ _____

Bankcard Visa Mastercard

Name on Card _____
 Card Number _____
 Expiry Date _____
 Signature _____

Please mail completed form with payment to: _____ Or Fax to: _____

Reply Paid 66
 AUUG Membership Secretary
 PO Box 366
 KENSINGTON NSW 2033

AUUG Inc
 (02) 8824 9522

Section D: MAILING LISTS

AUUG mailing lists are sometimes made available to vendors. Please indicate whether you wish your name to be included on these lists:

Yes No

Section E: AGREEMENT

I/We agree that this membership will be subject to rules and by-laws of AUUG as in force from time to time, and that this membership will run from time of joining/renewal until the end of the calendar or financial year.

I/We understand that I/we will receive two copies of the AUUG newsletter, and may send two representatives to AUUG sponsored events at member rates, though I/we will have only one vote in AUUG elections, and other ballots as required.

Signed: _____

Title: _____

Date: _____

AUUG Secretariat Use

Chq: bank _____ bsb _____

A/C: _____ # _____

Date: _____ \$ _____

Initial: _____ Date Processed: _____

Membership#: _____



UNIX® AND OPEN SYSTEMS USERS

Membership Application

AUUG Inc

PO Box 366, Kensington NSW 2033, Australia

Tel: (02) 8824 9511

Free Call: 1 800 625 655

Fax: (02) 8824 9522

email: auug@auug.org.au

ACN A00 166 36N (incorporated in Victoria)

<http://www.auug.org.au>

AUUG Inc is the Australian UNIX and Open Systems User Group, providing users with relevant and practical information, services and education through co-operation among users.

AUUG OFFERS SOMETHING FOR YOU!!

Education
Tutorials
Workshops

AUUGN
Technical Newsletter
AUUG's quarterly publication, keeping you up to date with the world of UNIX and open systems.

Events.....Events.....Events
• Annual Conference & Exhibition
• Overseas Speakers • Local Conferences
• Roadshows • Monthly Meetings

DISCOUNTS
to all AUUG events and education.
Reciprocal arrangements with overseas affiliates.
Discounts with various internet service providers, software, publications and more...!!

Connections
• Newsgroup
aus.org.auug

Application for Individual or Student Membership

Section A: PERSONAL DETAILS

Surname _____ First Name _____
 Title: _____ Position _____
 Organisation _____
 Address _____
 Suburb _____ State _____ Postcode _____
 Telephone: Business _____ Private _____
 Facsimile: _____ E-mail _____

Section B: MEMBERSHIP INFORMATION

Please indicate whether you require Student or Individual Membership by ticking the appropriate box.

RENEWAL/NEW INDIVIDUAL MEMBERSHIP

Renewal/New Membership of AUUG \$110.00

RENEWAL/NEW STUDENT MEMBERSHIP

Renewal/New Membership of AUUG \$27.50
 (Please complete Section C)

SURCHARGE FOR INTERNATIONAL AIR MAIL \$66.00

Rates valid as at 1 March 2000. Memberships valid through to 30 June 2001 and include 10% GST.

Section C: STUDENT MEMBER CERTIFICATION

For those applying for Student Membership, this section is required to be completed by a member of the academic staff.

I hereby certify that the applicant on this form is a full time student and that the following details are correct.

NAME OF STUDENT: _____

INSTITUTION: _____

STUDENT NUMBER: _____

SIGNED: _____

NAME: _____

TITLE: _____

DATE: _____

Section D: LOCAL CHAPTER PREFERENCE

By default your closest local chapter will receive a percentage of your membership fee in support of local activities. Should you choose to elect another chapter to be the recipient please specify here:

Section E: MAILING LISTS

AUUG mailing lists are sometimes made available to vendors. Please indicate whether you wish your name to be included on these lists:

Yes No

Section F: PAYMENT

Cheques to be made payable to AUUG Inc
 (Payment in Australian Dollars only)

For all overseas applications, a bank draft drawn on an Australian bank is required. Please do not send purchase orders.

-OR-

Please debit my credit card for A\$ _____

Bankcard Visa Mastercard

Name on Card _____

Card Number _____

Expiry Date _____

Signature _____

Please mail completed form with payment to: Or Fax to:

Reply Paid 66
 AUUG Membership Secretary
 PO Box 366
 KENSINGTON NSW 2033
 AUSTRALIA

AUUG Inc
 (02) 8824 9522

Section G: AGREEMENT

I agree that this membership will be subject to rules and by-laws of AUUG as in force from time to time, and that this membership will run from time of joining/renewal until the end of the calendar or financial year.

Signed: _____

Date: _____

AUUG Secretariat Use

Chq: bank _____ bsb _____

A/C: _____ # _____

Date: _____ \$ _____

Initial: _____ Date Processed: _____

Membership#: _____