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

Editorial

Con Zymaris
auugn@auug.org.au

As software and systems professionals, we should be well versed in the intellectual property and licencing issues which surround the technology we work and play with. One such occurrence which had some attention in recent months is that of the locally developed, highly advanced packet-filtering technology, IP Filter.

IPFilter was written in Australia by a long-time associate and former work colleague, Darren Reed. Darren is a very competent coder, and internationally noted security authority. He has presented papers at past AUUG Conferences (hey, *plug!* :-), and is an active member of the local Unix scene.

The reason I introduce IPFilter here is because of the problems which arose when many people who used this tool were recently hit by the reality of its licence. Many people thought that IPFilter, which has a solid following amongst the BSD Unix community and is included in platforms like OpenBSD, must be released under an open source licence. In fact, this wasn't the case. The furor erupted when Darren clarified the licencing issue, much to the chagrin of developers from groups within the OpenBSD community. Neither they nor Darren had done anything wrong, it was just a mismatch of expectations. If the users of IPFilter had fully read and understood the licence, rather than just presumed, this issue would have been avoided.

It is therefore important that we all understand what rights and responsibilities we have been allocated by the licence which the software we use is shipped out with. We also have to broaden our outlook and be vigilant against powerful vested-interest industry lobby groups which act to enhance the rights of vendors in the licenced intellectual property products (software, music, ebooks, movies) to the detriment of users.

Cheers,

Con

Contribution Deadlines for AUUGN in 2001

Volume 22 • Number 3 – September 2001: **August 30th, 2001**

Volume 22 • Number 4 – December 2001: **November 17th, 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

President's Column

David Purdue
David.Purdue@auug.org.au

> **remedy** /'remIdI/ -n. (pl. -ies) (often foll. by for, against) 1 medicine or treatment. 2 means of counteracting or removing anything undesirable. 3 redress; legal or other reparation. - The Pocket Oxford Dictionary

It would appear that, for the moment at least, Microsoft has been let off the hook. The US court of appeal has considered the antitrust case that the Department of Justice brought against Microsoft, and found that while Judge Thomas Jackson had reached the right conclusion (Microsoft has acted in a monopolistic fashion), the remedy he proposed (splitting Microsoft in two) was not appropriate.

Perhaps Judge Jackson's proposed break up of Microsoft was over the top - and the point made by the appeals court was that Jackson had not provided sufficient justification that splitting the company in two would solve the problems. But I don't think you could blame Judge Jackson if he was pissed off at Microsoft. From the reports of court activity it would seem that Microsoft treated him like an idiot and thought that the court proceedings had nothing to do with their future operations; they thought they were untouchable.

Which begs the question: what is the correct remedy to Microsoft's anticompetitive behaviour?

I don't think a simple split in two would have solved the whole problem. The proposal was to split Microsoft into the Operating System company and the Everything Else company. Why split this way? Well, OS accounts for about 50% of Microsoft's revenue and profit, so you end up with two companies about the same size. But would this split on its own stop either corporation from acting in an anticompetitive manner? The main problem is Microsoft using dominance in one market to achieve dominance in another, and while it is the OS that provides the most leverage, Everything Else still contains a lot of dominant market positions, and a lot of scope for cross promotion.

Even a split into three, namely OS, Applications and Everything Else, would not entirely solve this behaviour unless severe restrictions were placed on the communications these companies could have and the deals they could strike. This could end up being unenforceable.

But perhaps the Open Systems and Open Source community can teach the US judiciary a thing or two about competition. I invite those who do not think that competition reigns in the Open Source world to visit my office and listen to the proponents of Mandrake Linux and Debian Linux slug it out.

Competition is rife within the Open Systems world because it is IT that is based on open standards. For example, let's look at e-mail systems. There are a large number of different e-mail systems available, some commercial, some open source. How do you choose the one that's right for you? You look at cost, performance, the quality of the implementation, the level of support you can get from the vendor, and a multitude of other factors. However, one thing you take as a given in making this choice - your e-mail system will be able to talk to everyone else's e-mail system because it implements the SMTP standard. Any failure to adhere to this standard will be reported to the vendor of this system (or project group if it is open source) as a bug.

So perhaps that is the direction we should be taking with Microsoft. We force them to publish the API's used by their operating systems, we force them to publish the file formats used by their applications, and we force them to adhere to these standards. In addition, if they decide to change these API's or file formats, then the changes should be published when they are decided upon, which should be in advance of the release of the products that use these changes.

The aim of this is to flatten the playing field. If other companies (and open source projects) have access to the standards used in Microsoft software, then Microsoft will be forced to compete on the basis of quality of implementation, support and price.

/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: Fri, 20 Apr 2001 15:44:20 +1000
From: Grant Morphett <gmorph@canb.auug.org.au>
To: auugn@auug.org.au
Subject: Congrats

I usually read AUUGN over a coffee and find it quite good but the latest issue (Volume 22 Number 1) I thought was excellent. Just wanted to pass on a hearty well done to all involved.

cheers

--
Grant Morphett
GMORPH CONSULTANTS Pty Ltd -
Solutions Outside the Square
tel : +61 (0) 403 395486

Date: Mon, 18 Jun 2001 05:12:37 +1000
From: Greg Black <gjb@gbch.net>
To: talk@auug.org.au
Subject: Re: Perl and Unix

|> Is Perl to be considered part of a Unix installation
|> (i.e. can you presume that the base
|> libraries/modules will be on a Unix system) or is it
|> still "foreign"?
|
| Depends which unix you're talking about. For ex
| ample Solaris 8 has
| perl, but earlier versions of Solaris don't include
| perl.
|
| Also, you can install Linux, *BSD, Solaris 8, etc
| without their perl
| packages, so it is not a given that they will always be
| there.

You certainly can't install BSD/OS or FreeBSD without Perl; I'm not a user of the other BSD variants, so can't speak for them.

If only they'd start including Python by default, then we'd have a nice tool that we could depend upon instead of that Perl mess.

Date: Mon, 18 Jun 2001 15:21:48 +1000 (EST)
From: David Purdue <davidp@canb.auug.org.au>
To: Glenn Satchell <Glenn.Satchell@uniq.com.au>
cc: talk@auug.org.au, gjb@gbch.net
Subject: Re: Perl and Unix

On Mon, 18 Jun 2001, Glenn Satchell wrote:
> But typically the version of perl/python/whatever
> that comes with an OS release is some number of
> versions behind the latest and greatest so you prob
> ably want to build or install your own version any
> way.

Depends on what you want it for. If you are writing a whole bunch of perl cgi scripts that make use of the latest tweaks on a daily basis, then sure, build it.

For me - I just want the 'perl available' check to pass when installing other packages, so having perl arrive as part of the Solaris 8 install saves me the time it would take to download and build it. (Ditto OpenBSD, etc.)

DavidP

AUUG2001 Ad

AUUG2001 Ad

Membership Renewal Reminder

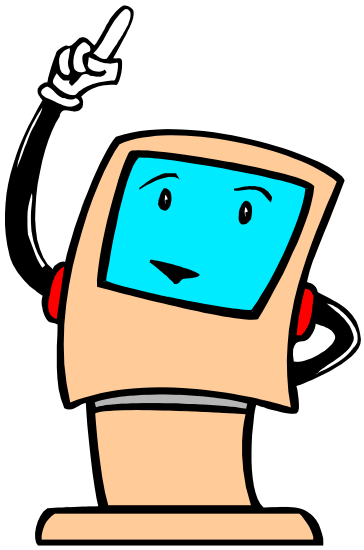

If...

your AUUG membership expired on
30th June 2001 and
you do not send in your membership renewal

This will be your LAST Copy of AUUGN

Membership Renewal invoices were mailed in May, however, if you have not yet received yours, please contact:

Liz Carroll, AUUG Business Manager
at busmgr@auug.org.au

or call the AUUG Office on
1-800-625 655

Public Notices

Upcoming Conferences

August 13-17
10th USENIX Security Symposium
Washington, D.C.

September 17-20, 2001
The O'Reilly Peer-to-Peer Conference
Omni Shoreham Hotel, Washington, DC

November 6-10
5th Annual Linux Showcase and Conference
Oakland, CA

November 8, 2001
XFree86 Technical Conference
Oakland Convention Ctr
Oakland, California

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 11-14, 2002
BSDCon 2002
Cathedral Hill Hotel
San Francisco, Californi

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

My Home Network (June 2001)

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

I've finally done it, I've set up a new network link, entirely separate from ANSTO. If you don't know, up until now, my sole connection to the Internet has been via my previous employer, ANSTO, hidden behind their firewall. Unfortunately, all good things must come to an end, and since I'm no longer an ANSTO employee, I could no longer continue using their connection.

Of course, as I had to change, I decided it had better be a good upgrade, so instead of changing to another 56K connection with some ISP, I moved to an ADSL connection with BigPond. This gave me almost a ten--fold increase in bandwidth, from 56K up/33K down (if you don't understand this, look up how 56K modems work) to 512K up/128K down.

These changes entailed considerable work to firstly establish a connection and secondly, to rework all my services to work with the new systems.

The first part of this, establishing a connection, involved some new software, and an understanding of how ADSL works. ADSL or Asymmetric Digital Subscriber Line is described by Telstra as a new "super--fast technology", with a speed up to 1.5Mbps, which runs over the normal copper phone line back to the Telstra exchange. This isn't really super-fast but compared to a modem it is an improvement. It is one of the new technologies, that are grouped as broadband and includes cable. The fact that it runs over normal lines is a huge advantage, and even better, it is possible to still use a normal phone or modem on the same line. The only requirement is the use of a filter, which is supplied by Telstra.

An application for an ADSL connection can be made directly from Telstra's web page, http://www.telstra.com.au/adsl, and provided it is available in your area, a Telstra technician will turn up with an "ADSL modem", software and connection details. Of course their software only runs on certain proprietary operating systems, and during the setup you should allow them to prove it works on such a system.

One important point is the "ADSL modem". This is not really a modem, in fact being more like a router than a modem and comes in two possible types, a UTP Ethernet connection and a USB connection. For an Ethernet connection, Telstra will also supply an Ethernet card. For Unix systems it is best to select the Ethernet connection.

Now, one final thing you need to know, Telstra like most ADSL providers run a strange protocol called PPP (Point-to-Point Protocol) over Ethernet (PPPoE). So rather than having a normal IP connection, ADSL takes PPP packets and stuffs them into raw Ethernet packets, directed to/from the network interface of the ADSL modem. This means that the interface does not need an IP address and, in fact, it is recommended not to configure one.

Okay, so while Telstra do not support ADSL connections to Unix/Linux systems, there is enough support elsewhere to more than make up for it. So the first thing to do is install an extra Ethernet card and get it working before starting. You can use the one supplied by Telstra, or some other cheap one, all you need is a 10Mb/sec connection. Even a suitable UTP comes along, so there is no need to stay up all night trying to wire one up.

Once the hardware is in place, the next item is software. While a quick search of freshmeat.net will turn up a few PPPoE clients, there appears to be one that is used by the majority of Unix users, and the only one that was recommended when I asked. The package you require is rp-pppoe, version 3.0, written by Roaring Penguin, available at http://www.roaringpenguin.com/pppoe. While this software is written for Linux, it also runs on Solaris and NetBSD.

The only other piece of software required is the PPP daemon, pppd, version 2.3.7 or later, this runs on Linux Kernel 2.2. If you are running the Linux 2.4 kernel, then you can enable kernel-mode support, however, I won't cover it in this article.

rp-pppoe is very simple to install, basically, either from an RPM (an old one is included in Red Hat distributions) or the source available on Roaring Penguins site. Once installed, it is easy to configure by running the script "adsl-setup". This asks a few questions for the setup: your BigPond user name and password (supplied by Telstra), the Ethernet interface (by default "eth1"), the DNS setup and probably the most misunderstood, the type of firewalling.

One of the first things I noticed after setting up my BigPond connection was the number of port scans continually running past my site. Unfortunately, at the present time the Internet is a wild place, and not something you can be unprotected on. Roaring Penguin have given three options: the installation of no rules, leaving it up to you to protect yourself; the installation of rules suitable for a stand-alone system; or the installation of rules suitable for an internal gateway.

All these rules make use of ipchains, which is the standard for Linux 2.2, but must be compiled into the kernel, and are seen as a basic starting point. In reality for all but the simplest case you will need to extend them, and I'll cover some of this later. As you should be able to guess from previous articles, my choice here was the internal gateway or "MASQUERADE".

Not to get into too much further discussion about security, as plenty can be found on the various BigPond news groups, but in the week prior to preparing this article my gateway machine was scanned at least 16 times. Obviously some people have nothing better to do, and there is too much free computing power around.

At this point things were looking good. Ignoring the time I'd spent prior to the installation planning, it had been about 4 hours since the Telstra technician had arrived to setup my ADSL connection, and in that time the physical installation had been done (about 1 hours, at a leisurely pace), installed an additional Ethernet card in my network server, installed and configured the Roaring Penguin software and was ready to go online with my entire network. So after crossing and uncrossing fingers, the next step was to run "adsl-start". If anything, this was an anti-climax, all that happened was the lights blinked a bit and I was online.

But now things started to go wrong.....

Well not that badly, and not too unexpectedly, but the next step was to get my system to use the ADSL link, and not my previous link to ANSTO. The first and most critical item was getting DNS working correctly, and this became obvious almost immediately, as everything I tried to use over the network failed. In fact there was a step before this, try and prove basic network connectivity, and initially it didn't seem to work. The two tools I used were 'ping' and 'traceroute', and as I didn't have any DNS, it was back to IP addresses. The first host to try was one of BigPond's DNS servers, which produced no response, next try the other, again nothing. At this point I switched back to the system I had setup with the Telstra engineer, and guess what, they couldn't ping the servers either. Anyway, to cut a long story short, I soon found out that the BigPond servers don't respond to ping, but there are plenty of other systems that do. To make it even more fun, traceroute from the server still fails to work, although as I've found out since, it does work from other hosts within my network.

Anyway, the easiest thing I found to do was to de-construct much of my network infrastructure and gradually bring each item back online. The first step was to stop 'named' and just use "/etc/resolv.conf". This worked, so next I reconfigured "/etc/named.conf" to use BigPond's DNS servers as forward-first servers, allowing it to go out directly as well.

In addition to DNS I found other problems with 'squid', 'ntp' and 'mirror', all different and all of which meant I had to go back to basics with each package.

For example, BigPond does not have a caching server, or else it is run in transparent mode. My previous squid configuration passed all requests to ANSTO's cache, but this was no longer possible.

With time synchronisation using 'ntp' it was a different problem, security blocked me. The servers I used before were now blocked by ANSTO's security, and something that took longer to notice, the default IPChains policy was stopping any responses getting back to the daemon.

To understand this, you need to understand a bit about the differences between TCP and UDP. TCP is a connection-oriented protocol, and each packet includes details to allow it to be identified as part of a connection. The initial packet in a connection is also specially marked (with a SYN bit) and if it does not get through no connection is established and any subsequent packets are just dropped. The fact that you can filter these SYN packets is used by many firewalls to handle TCP connections. In the default IPChains setup for rp-pppoe, locally originated connections are allowed, to send out SYN packets, but any external attempts to create a connection find that the initial SYN packet is just dropped. In addition for extra protection all ports in the privileged range (i.e. <1024) are automatically dropped. Most TCP based connection methods generally initiate a connection from a non-privileged port and so all works fine.

However, UDP is a connectionless protocol, so each packet stands alone. There is no way to tell if it is associated with any previous ones, and most UDP based connections are directed to a specific fixed port. For most well-known daemons, this is often in the privileged range. For example the NTP protocol expects all packets to be directed to UDP port 123.

The way around this problem was to configure a special rule that allowed connections to UDP port 123. This is an obvious security concern, and a reason for keeping up to date with relevant security patches. The hole this presents could be limited with a number of methods which I will not go into here.

Security was also the reason for problems with the 'mirror' program, which is used to automatically keep my patches up-to-date. However, this was a different problem. 'mirror' uses FTP to transfer the files, and as is widely known in the security community, FTP by default, allows a connection to be established from the remote site back to the local client. In terms of security, this is a problem, and while the rules for IPChains can be bent to accommodate this, it is better to use the PASV command, which allows all connections to be initiated by the local client. The 'mirror' script did not use PASV by default and it took some digging to find how to support it. In the end it only required the inclusion of the option "passive_ftp=true" in the default configuration file, "/etc/mirror.default".

I should mention that this isn't a very big issue. Most modern FTP servers will accept the PASV command, and most clients will automatically use it, most even have it as the default. In addition, for masqueraded clients, there is a special module that runs on the server to automatically handle the process, although this doesn't help if IPChains does not allow non-PASV FTP. Both these make it a non-issue for most FTP access. Unfortunately, it sometimes takes a while to notice problems. For example, it was a month before I realised that Symantec's LiveUpdate for Macintosh couldn't do PASV, and that is why it was failing. Oh well, back to manual downloads and updates.

So, after a couple of days work, all this setup had been completed and the system had a basic level of security. Next it was on to making the system usable, without compromising the security.

One of the first things here was getting squid working fully. Firstly, it was necessary to unravel the strange setup BigPond has for their network. The first item I found was that many of the URL's refer to hosts

without a domain name, e.g. "http://www/usage.html". As I had my own domain this was interpreted as "http://www.crawford.emu.id.au/usage.html", which was wrong. The quick easy fix for this was to add the directive in "/etc/squid/squid.conf" which said:

```
append_domain .vic.bigpond.net.au
```

so it would append the domain "vic.bigpond.net.au" to the end of any unqualified names. (At this point I should mention that at present all BigPond's ADSL customers are in the 'vic.bigpont.net.au' domain, no matter where you live in Australia.)

The next, and for some time, more perplexing problem was that I couldn't connect to a number of BigPond's internal servers, and I believed that it was due to some restrictions by BigPond. However, after some further investigation I discovered that the problem was that the default squid configuration blocked SSL access to ports other than 443 and 563, and BigPond uses SSL to a number of other ports. The easiest way to fix this was to add another directive of the form:

```
acl SSL_ports port 1025-65535
```

i.e. allow SSL access to all ports above 1024.

This was followed by a more detailed analysis of what ports I was exposing to the outside network and why. If you study the firewall configuration, you will see that it allows access to UDP ports greater than 1024, and there are in fact a number of these. Probably the most notable is NFS, as it is a UDP based protocol, and despite officially using the 'portmapper', it in fact is always located on port 2049. As well, most of the RPC programs should not be exposed, and are all on ports from 1024 and up.

The approach I took here was to inspect the output of "netstat -an" looking for UDP ports greater than 1024 and specifying a specific IPChain deny command for that port. This is not the best method, as it does not automatically cater for future changes. One of the basic tenets of security is to deny everything that is not explicitly allowed, yet in this case I am trying to explicitly deny ports. Over time I will look to improve this.

At the same time while trying to block ports, I also went through and disabled services on the external interfaces that were not needed. For example, by default "squid" is available on all interfaces, as well as enabling an ICP and SNMP UDP port, again on all interfaces. To counter this I turned off SNMP and ICP, and bound the normal proxy port only to the internal interfaces, with directives of the form:

```
http_port bits.crawford.emu.id.au:3128
http_port localhost:3128
icp_port 0
snmp_port 0
```

Similarly, with "named" it was important to add the "listen-on" directive in the "options" block of "/etc/named.conf":

```
listen-on { 127.0.0.1; 203.16.204.5; };
```

You should however remember that there will be some ports available from the external interface, such as by "named" to send and receive transfers to other DNS servers. In general these can be either a fixed or random port, depending on the configuration directives.

So that brought the system to one I could use, with most services I wanted, but with one notable exception, mail. Because I didn't have a fixed IP address, there was no way I could get mail delivered to me, or was there? Yes, there was, by the use of a Dynamic DNS service, which is a story of its own, and one I'll cover in my next column.

For those of you who haven't noticed, I've again agreed to be Programme Chair for the AUUG Conference, and elsewhere in this edition of AUUGN you will be able to read the status of preparations. I hope to see you in Sydney for this conference.

National Linux Installfest: 2001

AUUGN would like to draw your attention to the preparations presently underway for this years National Installfest. The Installfests have become a major part of the yearly Linux event calendar, often attracting several hundred participants; many of them totally new to the ways of Unix. As such, we feel that events like this, which get national IT media coverage, are a great way to spread the word about the qualities and ethos of the Unix platform. If you have the skills, the enthusiasm or both, please sign up to support your State's LUG.

Details

from: Sarah Bolderoff <sarah@cs.unisa.edu.au>

The date is the 25th of August 2001.

The mailing list for national scale organisation: installfest-org@auug.org.au It's a good idea for local user groups to have their own mailing list for local organisation.

The web site will be at installfest.linux.org.au

AUUG is willing to provide support for user groups that wish to be involved in the installfest but aren't incorporated and don't have insurance.

For information on last years installfest go to: www.linux.org.au/installfest

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

I would like to put together a list of participating lugs, so people/lugs interested in joining in the installfest fun, can email me, sarah@cs.unisa.edu.au

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

Installing mod_gzip with Apache and PHP

Author: Donncha O Caoimh <donncha.ocaoimh@tradesignals.com>

Introduction

Mod_gzip, at

http://www.remotecommunications.com/apache/mod_gzip/

is a module for Apache that allows you to compress content from an Apache web server on-the-fly. It uses the same compression as gzip and no plugins or extra software is needed by your browser to take advantage of this product. Reduction in size of up to 90% or more is possible.

Install

1. Binary install
Download the binary for your platform and follow the instructions in the README under "HOW TO INSTALL"

2. Source install
Download mod_gzip.c and place it in the Apache source directory, ie $HOME/apache_1.3.XX/

If you're using v 1.3.17.1a of mod_gzip then apply the following patch by saving it to the same directory as mod_gzip.c: mod_gzip.1.3.17.1a.patch and running the following commands:

```
patch mod_gzip.c < mod_gzip.1.3.17.1a.patch
```

This solves a problem fixed by Artem Koutchine. Compile it as a DSO or statically into the web server itself. This is documented in the Apache README.configure ie. "--add-module=mod_gzip.c"

Install Apache with "make install"

Configuration

Add the following to httpd.conf

```
# MOD_GZIP configuration
mod_gzip_on Yes
mod_gzip_minimum_file_size  1002
mod_gzip_maximum_file_size  0
mod_gzip_maximum_inmem_size 60000
mod_gzip_item_include mime "application/x-httpd--
php"
mod_gzip_item_include mime text/*
mod_gzip_item_include mime "httpd/unix-directory"
mod_gzip_dechunk Yes
mod_gzip_temp_dir "/tmp"
mod_gzip_keep_workfiles No
mod_gzip_item_include file "\.php3$"
mod_gzip_item_include file "\.txt$"
mod_gzip_item_include file "\.html$"
mod_gzip_item_exclude file "\.css$"
mod_gzip_item_exclude file "\.js$"
```

and save the file.

Run

```
../bin/apachectl configtest
```

to make sure everything is ok, and then restart the webserver with

```
../bin/apachectl restart
```

Testing

It's never a good idea to introduce new software without testing. Here's some thoughts on how to test it without distrupting your users browsing (too much).
Run the web server on a different port. Try port 8080 and hit your webserver with the URL http://yourhost:8080/ in as many browsers as you can find.

Surround the above configuration with "Directory" directives so that mod_gzip is only active in one directory, and not site-wide.

```
<Directory
"/usr/local/apache/htdocs/testingdirectory/">
MOD_GZIP configuration
</Directory>
```

Use the Logformats described in the README to find out what is and what"s not being compressed.

Test all your applications to verify any device that your site supports will play nicely with mod_gzip.

Conclusion

Mod_gzip is amazing. You'll find large 100K HTML pages come down the wire as 12K compressed files. Your users will be very impressed by the new speed of your site.

From my viewing of several log files, it looks like MSIE is fairly hit and miss as to whether it will compress data or not. Netscape in Linux works very well, and the KDE Konqueror browser supports compression too.

It appears that Java applets don't support compression, even though the browser they're running in might do so. (Only Netscape tested)

If you use mod_gzip, subscribe to the mod_gzip mailing list, as any web content on mod_gzip grows out of date fairly quickly (including this one I guess).

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

http://www.linux.ie/articles/tutorials/mod_gzip.php

Joining the 6bone

By: Michael Paddon <michael@paddon.org>

Why Is IPv6 Interesting?

By now, you've probably heard of the next generation Internet Protocol, IPv6. While it provides many improvements and new capabilities, the driving force behind it's adoption is likely to be the much larger (and more flexible) address space that it defines. Continuing growth in the population of IP enabled devices has already put severe stress on address allocation and the routing infrastructure, and the roll out of new enabling technologies such as 3G wireless and broadband to the home will predictably create a new wave of demand.

One way of dealing with these pressures is to use address translation technologies and accept the consequential degeneration and balkanisation of global connectivity. Another path is transition to a networking technology that can support the demands of today and tomorrow. In discussing this matter with my peers, I am often surprised by an ingrained reluctance to change, usually conjoined with an assertion that NAT solves the addressing problem. To me this is akin to defending the DOS 640K limit, and shows a marked lack of imagination about how we might be using the net in ten or twenty years' time.

Where Do I Get IPv6 Software?

The first step to using IPv6 is, of course, finding a suitable implementation. The good news is that, if you are using open source, the chances are that you are already IPv6 ready. Current versions of {Free, Net, Open}BSD and Linux have a working version 6 stack out of the box. Solaris, AIX and MacOS X are reportedly in the same situation, whilst users of other flavours of Unix should check with their vendors. If you want to run IPv6 on your router appliance, you'll find that most major brands have the code either in beta test or ready for production.

IPv6 stacks are designed to cohabit peacefully with IPv4 (the "classic" version), and a router or host can use both concurrently. This is the first step in the migration to the new standard... there will never be a "flag day" when everyone switches to IPv6; rather we will see a gradual increase in version 6 traffic over a number of years until it is the dominant form of internet datagram.

For instance, by using *ping6(8)* on my OpenBSD box, I can see that I am already running IPv6 without having gone to any special effort:

```
$ ping6 -c 5 localhost
PING6(56=40+8+8 bytes) ::1 --> ::1
16 bytes from ::1, icmp_seq=0 hlim=64 time=0.16 ms
16 bytes from ::1, icmp_seq=1 hlim=64 time=0.144
ms
16 bytes from ::1, icmp_seq=2 hlim=64 time=0.136
ms
16 bytes from ::1, icmp_seq=3 hlim=64 time=0.14 ms
16 bytes from ::1, icmp_seq=4 hlim=64 time=0.131
ms

--- ::1 ping6 statistics ---
5 packets transmitted, 5 packets received, 0%
packet loss
round-trip min/avg/max/std-dev =
0.131/0.142/0.160/0.010 ms
```

All the examples in this article were produced using OpenBSD 2.9. The commands should be pretty much the same for all the BSDs (although your interface names and addresses will differ), whilst other Unix variants may require slightly different invocations. Check your local man(1) pages for more details.

Using IPv6 on a LAN

The next obvious step is to get two machines on your LAN communicating using IPv6. The machine on my desk, called lum, is equipped with a "rl" network card with the following configuration:

```
lum$ ifconfig rl0
rl0: flags=8843 mtu 1500
  media: Ethernet autoselect (none)
  status: active
  inet 203.25.251.2 netmask 0xffffff0 broadcast
203.25.251.15
  inet6 fe80::260:67ff:fe06:19a5%rl0 prefixlen 64
scopeid 0x1
```

You can see that there is an IPv6 address configured for the interface, although it looks very different, and perhaps a little awkward, compared to a classic quad number address. Don't worry... it will all start to make sense soon enough, and remember that in the real world we type in DNS names and not addresses anyway.

This address was automatically generated and configured when I enabled the network interface card, based on its underlying MAC address (yes, there is a standard for this). It is called a "link-local" address and is only valid on the network that the interface is directly connected to. The ability to auto-configure addresses, without engaging in a protocol such as RARP, is a powerful and robust bootstrap mechanism.

An IPv6 address is 128 bits long (that's right... no more running out of address space until we go extra-terrestrial), and is conventionally written as 8 lots of 16 bit hex numbers, separated by colons. You are allowed to have one instance of two consecutive colons to show that all the digits in between are zeroes... For instance, "::1", the loopback address, is the same as "0:0:0:0:0:0:0:1" but a lot easier to type.

IPv6 address come in many flavours, which can be distinguished by the first few bits of the address. A few common
examples are:

Prefix	Address type
0000 0000	Reserved
001	Aggregatable Global Unicast
1111 1110 10	Link-Local
1111 1110 11	Site-Local
1111 1111	Multicast

Since link-local addresses are only good for a particular network, you have to specify the network interface as well. Conventionally, this is done with a trailing

percent sign plus the interface name. So now the address "fe80::260:67ff:fe06:19a5%rl0" actually makes sense.

I have another machine, called gunbuster, with a "de" card on the same LAN. Let's check out the interface configuration:

```
gunbuster$ ifconfig de0
de0: flags=8863 mtu 1500
        media: Ethernet autoselect (10baseT)
        status: active
        inet 203.25.251.1 netmask 0xffffff0
broadcast 203.25.251.15
        inet6 fe80::200:c0ff:fe4b:fdb%de0 prefix-
len 64 scopeid 0x1
```

I can ping gunbuster from lum, using the existing link-local addresses. Note that I have to remember to tack lum's network card name to the end to gunbuster's link-local address, or I'll get a "no route to host" error.

```
lum$ ping6 -c 5 fe80::200:c0ff:fe4b:fdb%rl0
PING6(56=40+8+8 bytes)
fe80::260:67ff:fe06:19a5%rl0
     --> fe80::200:c0ff:fe4b:fdb%rl0

16 bytes from fe80::200:c0ff:fe4b:fdb%rl0,
icmp_seq=0 hlim=64 time=0.445 ms
16 bytes from fe80::200:c0ff:fe4b:fdb%rl0,
icmp_seq=1 hlim=64 time=0.44 ms
16 bytes from fe80::200:c0ff:fe4b:fdb%rl0,
icmp_seq=2 hlim=64 time=0.437 ms
16 bytes from fe80::200:c0ff:fe4b:fdb%rl0,
icmp_seq=3 hlim=64 time=0.441 ms
16 bytes from fe80::200:c0ff:fe4b:fdb%rl0,
icmp_seq=4 hlim=64 time=0.443 ms

-- fe80::200:c0ff:fe4b:fdb%rl0 ping6 statistics --
5 packets transmitted, 5 packets received, 0%
packet loss round-trip min/avg/max/std-dev =
0.437/0.441/0.445/0.003 ms
```

Sitewide Addressing

No one in their right mind would use a link-local address for anything but low level configuration and diagnosis.

If you are using IPv6 solely within a given organisation, the next step you can take is to assign site wide addresses to your nodes. As with IPv4, you can technically assign any addresses you like, but you would be smart to stick to the defined "site-local" scheme and ensure future interoperability. Site-local addresses are rather like the RFC-1918 addresses used internally by many organisations today, in that they are guaranteed to never clash with a globally routable address.

Site-local addresses are currently specified to be of the form "fec0:0:0:S:I:I:I:I", where "S" is the 16 bit subnet identifier, and "I:I:I:I" is the 64 bit node id. You can assign subnet and node identifiers as you see fit. There is a defined algorithm, for instance, for turning a 48 bit ethernet MAC address into a node id (which is the same one used for creating link-local addresses). However, I prefer to assign addresses manually, because I have only a few machines and I don't like addresses to change when I swap interface cards.

I can assign site-local addresses to my machines with a few simple commands:

```
gunbuster$ ifconfig de0 inet6 alias fec0::1
lum$ ifconfig rl0 inet6 alias fec0::2
lum$ ping6 -c 5 fec0::1
PING6(56=40+8+8 bytes) fec0::2 --> fec0::1

16 bytes from fec0::1, icmp_seq=0 hlim=64
time=0.557 ms
16 bytes from fec0::1, icmp_seq=1 hlim=64
time=0.444 ms
16 bytes from fec0::1, icmp_seq=2 hlim=64
time=0.439 ms
16 bytes from fec0::1, icmp_seq=3 hlim=64
time=0.429 ms
16 bytes from fec0::1, icmp_seq=4 hlim=64
time=0.421 ms

--- fec0::1 ping6 statistics ---
5 packets transmitted, 5 packets received, 0%
packet loss round-trip min/avg/max/std-dev =
0.421/0.458/0.557/0.050 ms
```

When you get sick of typing in raw addresses, you can map them to names in my /etc/hosts file or my DNS service.

Modern implementations of named support the "AAAA" record type, which is used just like an "A" record only it holds an IPv6 address. Talk to your friendly DNS administrator for more information.

Global Addressing

Before you can connect to other IPv6 sites, you need to understand the structure of "aggregatable global unicast" addresses, which are the types of addresses being allocated for wide area interconnectivity. It should be noted, however, that only 1/8 of the total possible address space is reserved for this format, leaving room for the future development of alternative global addressing schemes.

One problem that a huge address space does not address (and in fact could exacerbate) is that of routing table growth. Without some kind of topological structure, you potentially need a route for every destination in the universe which clearly does not scale well. As you might expect, aggregatable global unicast addresses deal with this problem by aggregating routes, in much the same way as is done today with the CIDR mechanism. Hence, addresses have a well defined structure:

Format Prefix	TLA ID	Reserved	NLA ID	SLA ID	Interface ID
001	<13 bits>	<8 bits>	<24 bits>	<16 bits>	<64 bits>

What are all these TLA's? Addresses are organised in a four tiered topological hierarchy:

1. Top level aggregator (TLA) identifiers are assigned to organisations at the top of the routing tree, typically major carriers and exchanges. There is currently room for 8192 such entities.
2. Next level aggregator (NLA) identifers are assigned by each TLA to create a routing architecture appropriate to their circumstances.
3. Site level aggregator (SLA) identifers are assigned by each end user organisation to create their private routing architecture.
4. Network interface identifiers are assigned to each host.

The TLA and NLA identifers are together referred to as the "public topology", since they define the collection of organisations which provide public transit services. The reserved bits between the TLA and NLA fields are there to allow for expansion of either or both in the future as demand and technical capacity dictates.

There may be additional structure imposed at the NLA and SLA level, created by the way that the identifiers are handed out. For instance, if a regional registry is given a contiguous block of NLA identifiers, it can break it up into sub-blocks for different carriers, who may in turn break their blocks up into smaller allocations again.

However, the Interface ID level may not be broken up in this way as the various standards for generating these identifiers demand a 64 bit space.

To date, three TLA assignments have been made, each for a different purpose:

Prefix	Use
3ffe::/16	6bone experimental testbed allocation
2001::/16	Regional Internet Registry production allocation
2002::/16	6to4 transitional address space

In practice, each of these TLAs structures their subordinate address space differently:

- 6bone Addresses The 6bone is a collaborative testbed for the deployment of IPv6, created so that early adopters can test implementations and gain operational experience. Each backbone site is assigned a "pseudo TLA" (pTLA) identifier, which hands out "pseudo NLA" (pNLA) identifers to connecting organisations, each according to its own policies.

Prefix and TLA ID	pTLA ID	pNLA ID	SLA ID	Interface ID
3ffe	<12 bits>	<20 bits>	<16 bits>	<64 bits>

  Historically, the pTLA field used to be 8 bits. Hence, as a special case, any pTLA identifier less than 0x0800 is treated as an 8 bit value, with the corresponding pNLA field lengthened to 24 bits. This piece of architectural ugliness would be bothersome except for the fact that the 6bone is supposed to be experimental, and the address space will be eventually reclaimed for other purposes. In practice, it has no impact unless you are a pTLA, and then it only complicates your routing tables a bit.

Production Addresses The three Regional Internet Registries (APNIC, ARIN and RIPE NCC) have begun allocating production IPv6 unicast address blocks. Since the lifespan of the 6bone is explicitly finite (although the termination date is currently undefined), it would be wise to begin considering production addresses for mission critical purposes. This will save the effort and expense of renumbering at a later date, although with IPv6, this is potentially much easier than with IPv4. Production addresses are organised along similar lines to the 6bone space, although there has been some fine tuning. In addition, while the concept of a pseudo TLA has remained, it has been re-dubbed a "subTLA".

Prefix & TLA ID	subTLA ID	Reserved	NLA ID	SLA ID	Interface ID
2001	<13 bits>	<6 bits>	<13 bits>	<16 bits>	<64 bits>

The only problem is that, unless you are a major carrier, exchange or ISP, you probably won't qualify for a subTLA identifier of your very own. Which means that you have to wait for your ISP to get one, and for them to allocate part of their space to you. Or maybe for your ISP's ISP to get one, etc, etc. Now while this may, on the surface, suck somewhat, remember that this sort of disciplined hierarchy will keep the routing tables small enough to keep working. And that is a good thing for us all. Meanwhile, if your ISP is trying to redefine "customer service" to mean "something that moves at the speed of flowing glass", then vote with your dollars. Failing that, join the 6bone until your part of the world catches up with the twenty first century.

- 6to4 Transitional Addresses Anyone who already has a globally unique IPv4 address also has officially assigned IPv6 address space via the "6to4" mechanism. These addresses are used as a means of connecting isolated IPv4 subnets into the IPv6 fabric, and is intended as a transitional mechanism. These addresses are of the form:

Prefix and TLA ID	IPv4 address	SLA ID	Interface ID
2002	<32 bits>	<16 bits>	<64 bits>

  In the 6to4 scheme, when I want to send a version 6 datagram to a 2002::/16 prefix, I simply encapsulate it in a version 4 packet and send it to the IPv4 address embedded in the address. Clearly, this requires appropriate encapsulation software and routing at all 6to4 sites, plus sites who are willing to relay traffic between 6to4 and non-6to4 sites. For those wishing to play with 6to4, the open source IPv6 stacks provide appropriate support out of the box.

Global Connectivity

Don't you ever wonder who bought the first telephone? Lord knows, that salesman deserved his commission! Early adopters of IPv6 face a similar problem. How do you build critical mass without a backbone, and how do you build a backbone without critical mass?

The answer is to build a virtual backbone, tunnelling IPv6 datagrams encapsulated in IPv4 packets. In other words, when two version 6 nodes wish to exchange traffic, they actually send version 4 packets whose payload is the IPv6 datagram. Each tunnel is a point to point link, and operates essentially as a physical connection.

With this technology, it is possible to establish a globally operational IPv6 network immediately, with native IPv6 physical links being established as traffic volume and performance demands.

The 6bone was established in exactly this manner, with the explicit goal of being a testbed for experimenters and early adopters. Right now, this is the best way to achieve global version 6 interconnection. Soon (hopefully), your ISP will be offering production address space along with supporting tunnels or physical links.

Alternatively, you may like to experiment with the 6to4 mechanisms, although it is early days yet in that domain.

Joining the 6bone is quite straightforward:

- Find an existing 6bone site to give you address space and connectivity. Usually this is a pTLA site, although it could just as easily be a pNLA site willing to slice off some of their SLA space. The current list of pTLAs can be found at the 6bone web site http://www.6bone.net/.

  You should try and choose a neighbour that is close (in IPv4 terms) to you. For instance, the only pTLA operating in Australia at the time of writing was Trumpet Software, which makes the choice easy for people in that part of the world.

- Set up the tunnel. On OpenBSD this is achieved using the gif(4) tunnelling interface:

```
$ ifconfig gif0 giftunnel
```

  On my machine gunbuster, my tunnel configuration looks like this:

```
gunbuster$ ifconfig gif0
gif0: flags=8051 mtu 1280
      physical address inet 203.25.253.124 -->
203.25.253.62
        inet6 fe80::200:c0ff:fe4b:fdb%gif0 -> ::
prefixlen 64 scopeid 0xe
```

  As you can see, the link comes up with an appropriate link-local address. It doesn't know what the address of
  the other end is, so it just shows it as the wildcard address "::" (all zeroes). There is no need to ever assign global unicast addresses to either end of the link since the endpoints do not need to be globally addressable.

- Set up the routing. Since we are an end node, we just point our default route at the tunnel:

```
gunbuster$ route add -inet6 default fe80::%gif0
```

  How does this work? When the tunnel is created, a route for fe80::%gif0 is auto-magically installed pointing to the virtual link (you can query all of your routes via netstat(1)). This saves us from ever having to know or rely on the link-local address at our peer's site.

- Define your internal address allocation and management policy, and assign global unicast addresses as appropriate. IPv6 addresses are configured on network interfaces in the usual manner:

```
gunbuster$ ifconfig de0 inet6 alias
```

- Test the link by contacting an external site:

```
gunbuster$ ping6 -c 5 gw.ip6.trumpet.net
 PING6(56=40+8+8 bytes) 3ffe:8001:15:101::1 -->
3ffe:8000:1::1

16 bytes from 3ffe:8000:1::1, icmp_seq=0 hlim=254
time=254.175 ms
16 bytes from 3ffe:8000:1::1, icmp_seq=1 hlim=254
time=163.672 ms
16 bytes from 3ffe:8000:1::1, icmp_seq=2 hlim=254
time=162.032 ms
16 bytes from 3ffe:8000:1::1, icmp_seq=3 hlim=254
time=144.572 ms
16 bytes from 3ffe:8000:1::1, icmp_seq=4 hlim=254
time=163.788 ms

--- gw.ip6.trumpet.net ping6 statistics ---
5 packets transmitted, 5 packets received, 0%
packet loss round-trip min/avg/max/std-dev =
144.572/177.648/254.175/38.940 ms
```

You can automate the tunnel configuration to occur at boot time by creating a /etc/hostname.gif0 file, containing something like this:

```
giftunnel 203.25.253.124 203.25.253.62
!/sbin/route add -inet6 default fe80::%gif0
```

Don't forget to also add alias lines to your other interfaces' configuration files, so they are assigned the global unicast addresses that you choose.

Now What?

When you have a working connection to the wider IPv6 network, the obvious question is: what can I use this for? I've already introduced ping6. As you might expect, there is a corresponding traceroute6, which will allow you to explore the topology of the network.

On an OpenBSD box, you will find that many network applications are already IPv6 enabled, including ftp(1), telnet(1) and ssh(1). All of these applications will use IPv6 if a version 6 address is specified or if the target has a "AAAA" DNS record.

Any server that is designed to work with inetd may be made available on the IPv6 network simply by adding an appropriate "tcp6" or "udp6" entry to /etc/inetd.conf. For instance, you can offer IPv6 time services by adding the line:

```
time stream tcp6 nowait root internal
```

Some IP stacks will listen for both version 4 and version 6 connections if a wildcard is specified. OpenBSD does not work like this, due to security concerns. As a consequence, a standalone server must be modified to listen on an IPv6 socket as well as (or instead of) an IPv4 socket. Patches for popular servers are becoming common.

Security

As you begin to experiment with IPv6, it is important to watch out for security holes. Many servers have code that applies some form of address based security, and this may break in unexpected ways when handed an IPv6 socket. For instance, my anti-spamming code,

which checks the incoming source address, failed miserably when I IPv6 enabled my SMTP server.

Similarly, many packet filters and security proxies are not IPv6 ready, and you may be creating ways to simply bypass your security infrastructure if you open up an IPv6 tunnel without sufficient forethought.

Last Words

There is an enormous amount of information that could not be included in this article for reasons of space and time. The intent of this document, however, is to give you enough practical advice to bootstrap yourself onto the the next generation Internet. From there, although there is an enormous and growing body of new facts to learn and technology to understand, you should be able to proceed in a steady and incremental fashion.

Did you ever wish that you were around when the Internet got started? This is a chance to be part of the complete re-engineering of the net, and its metamorphosis into something even larger and more ubiquitous than the original. Don't miss this opportunity, because the next one will be a long time coming...

References

- 6bone home page: http://www.6bone.net/
- IETF draft-ietf-ngtrans-6bone-ptla-00: 6BONE pTLA and pNLA Formats
- IETF draft-ietf-ngtrans-6to4-03: Connection of IPv6 Domains via IPv4 Clouds without Explicit Tunnels
- OpenBSD home page: http://www.openbsd.org/
- Provisional IPv6 Assignment and Allocation Policy: http://www.apnic.org/drafts/ipv6/ipv6-policy-280599.html
- RFC 2373: IP Version 6 Addressing Architecture
- RFC 2460: Internet Protocol, Version 6 (IPv6) Specification
- RFC 2722: 6bone Backbone Routing Guidelines

The OpenBSD CD-ROMs

By: Greg Lehey <Greg.Lehey@auug.org.au>

This issue of AUUGN includes CDs with OpenBSD version 2.9. Our recent reader survey showed that this is one of the most popular software packages, and we hope that by distributing it in this manner, we can gain mindshare for OpenBSD.

OpenBSD is one of the three free implementations of BSD UNIX. Compared to the others (NetBSD and FreeBSD), it places greater emphasis on code purity and security, as witnessed by the motto "Four years without a remote hole in the default install!". As a result, it's a favourite with security-conscious people. Take a look at http://www.OpenBSD.org/ for more information.

Unlike the other BSD projects, the OpenBSD project has no corporate sponsors. Its main source of income is the sale of CD-ROMs.

So why are we giving them away? That's a good question, and we had quite a discussion about it with OpenBSD's "benevolent dictator", Theo de Raadt. We argued that by distributing OpenBSD to the AUUG, we would make it better known to people who otherwise would never have thought of trying it. It's clear, of course, that a number of AUUG members did intend to buy this CD, and now they're getting it for free, thereby losing income to the OpenBSD project.

What can we do about it? Well, the AUUG doesn't have a financial interest, but we strongly recommend to you that if you use the CDs, please make a donation to the OpenBSD project. See http://www.openbsd.org/donations.html for further details. We suggest a minimum of US $18, which is what OpenBSD gets from the sales of the CDs. If you prefer, you can submit the donation via the AUUG: contact our business manager, Liz Carroll, at 1-800--625-655 to pay by credit card, or send cheques to AUUG Inc., PO Box 366,Kensington NSW 2033. AUUG will forward the complete sum to the OpenBSD project.

Two issues ago we had problems with the CD-ROM distribution. If by any chance your issue arrives without the CDs, please contact Liz Carroll (phone number above) for details.

My Home Network: The Rejoinder

Andrew McRae (amcrae@cisco.com)

After reading Frank's regular discussion of his home network, I have been telling myself that I should make my contribution and talk about my own home network. Being the great procrastinator I am, I managed to avoid this for months whilst telling myself (usually on aircraft flying to various parts of the world) "Yeah, I'll do that when I get home". So, here it is...

A home network to me is more than just something I play around with; I am a telecommuter, which means my home network is more than just a hobby, it's my _livelihood_. In spite of that, I do manage to balance off the need for a reliable, stable system with some fun and experimentation (all in the name of Getting Things Done, of course :-).

The story starts many years ago, of course, with a 1200 baud modem and a dumb terminal, along with complaints from my wife that she could never use the phone any more, since I had that computer thing on it. At this time (and I am talking about _years_ ago, like 13 or 14), I had cajoled and badgered my employer into setting up a phone line dedicated to a 1200 baud modem that was used for ACSnet access (who remembers ACSnet??), and since I seemed to be the default system administrator, I was able to sneak a dumb terminal and modem home with me. This was great! E-mail, fetchfile, and this Usenet thing! How times have changed... Anyway, everybody is probably sick of hearing How Times Have Changed (especially when you get a bunch of the BOFs(*) around a table at an AUUG conference dinner). It wasn't long before I had several modems at work, and some that ran at 9600 baud, which actually made a dial-up line useful. Eventually I had a PC running some variant of 386/BSD, which morphed into FreeBSD. By this time, due to domestic pressure, I had installed a second phone line, and I had a REAL network, even though it was only a SL/IP line to an Annex terminal server.

[*BOFs - Boring Old Farts]

Around this time, I moved house; even better, I was building a house (well, not me personally, I am a software guy, and I try to avoid hardware if I can), and just prior to this I had been crawling through the ceilings in the office laying Cat5 networking cable. So I figured that it may be a neat idea to do the same thing with my new house - as we all know, it's much easier to lay cables with the skin off the walls, and I didn't want to contemplate crawling through the wall space in a residential two storey house. So I figured I would lay Cat5 cable to two bedrooms (the children's rooms), the garage (hey, I needed _somewhere_ to put this old Sun3) and have the cabling wired back to my study where a hub would sit. So now I had a Real Network, even though at the time, the cabling was sitting in the walls and I hadn't actually connected it to a wall socket or anything, let along put a hub in. No matter, I had my second phone line, my PC running FreeBSD, this thing called the 'Web' was just starting to appear, and I was a happy camper... until employment underwent upheaval, and within 4 months of moving into my home network paradise, I was on a plane to San Jose, working for the Biggest Networking Company In The World (the acronym spells 'cisco').

Well, after 6 months of living in Silicon Valley, and another 4 months shuttling back and forth, I finally got a chance to spend some time fiddling and set up my home environment. As I imagine has happened to most of us, I have been through several generations of PC technology - 385, 486, Pentium, Pentium II and so on. And like most computer nerds, I had cobbled together various PC systems, upgraded, used the castoffs for other systems, sworn at flaky hardware, thought _terrible_ thoughts about Bill Gates, and generally barked my knuckles playing around with lots of different PC boxes. My castoffs, of course, ended up first in my son's bedroom, and secondary castoffs in my (younger) daughter's bedroom (if for no other reason, than I figured I did not want to be outnumbered as the only nerd in the family). What helped a lot, before the Aussie dollar became the Pacific Peso, was to take a couple of hours out of every trip back to cisco HQ to cruise through that wonderful geek paradise in California called Frys, and collect marvellously cheap bits of PC hardware. If you are a geek, and you ever go to California, you absolutely, positively MUST undergo the pilgrimage to Frys. Getting back to the story, at this stage I had collected a number of reasonable PCs, a couple of which are running Windows, being used as a games loader, one which was permanantly running FreeBSD, and another Windows box which only got turned on when some impolite marketing person sent me a Word document.

In terms of Net connectivity, I have suffered from living on the edge of Sydney. It was only a couple of years ago that the local exchange could run ISDN, so for some time I put up with a dial-up modem connection to the local cisco office. Lo and behold, ISDN was finally available, and I was one of the first users in my area to get connected; the router (cisco, of course) handled all the gory ISDN dial-up side of things, and just had a 4 port Ethernet hub to connect to. I even had my own 3 bit IP subnet, so I felt somewhat special...

Anyways, I figured it was time to _really_ build a Home Network, so I got the hammer out, broke open a couple of bits of wall plaster, found the Cat5 cable I had installed 4 years earlier, and wired up the wall sockets. I upgraded all the PCs to Windows 95, installed Ethernet cards, and bingo, we were a wired household. Well, so what if all my daughter wanted to do was to surf www.barbie.com...

So, what exactly do I have running now? For my main machine, I use a dual processor PC running FreeBSD (of course). This machine acts as a file and print server (running Samba, naturally) for all the other PCs. It has a fair whack of disc space, all running on Ultra Wide SCSI (thank you Mr Fry for those cheap controllers and discs). It also acts as a DNS server for the network, because at one stage the ever-wise cisco IT people took away my IP subnet and put in a router that did NAT. Whilst I may have felt joy at having a

Class net of 10.0.0.0 all to my very self, the reality was that my local hosts disappeared off the world radar screen, making it hard to do the stuff I wanted to do for my job. Running DNS on my FreeBSD box solved one problem. Recently some amount of sanity returned and there is a single company NAT gateway so that at least within cisco I have a unique 6 bit subnet (for what I want to do, that's fine).

I have resisted what Frank has done, and avoided making the box a squid server. I just shudder, thinking of all those Barbie web pages being stored on my sleek Ultra Wide SCSI drives...

A couple of years ago my boss forced on me a laptop, figuring that there was all that spare time I had on airplanes. So, along with the separate Windows box I kept for all those Office documents I kept being sent, we have more functioning, operating computers at my house then people, 5 in total (computers, not people). This is not counting the other boxes that occasionally get an airing, or transient systems that happen to get put together. It makes for a great environment to play networked computer games.

The network currently runs off a 10Mbit ethernet hub. Recently I said to myself, "Do I work for the world's biggest networking company or not? [urk, let's not look at today's stock price] Well, why doesn't my network run 100 Mbits? And what about this Gigabit Ethernet thingy?". Well, I figured that perhaps Gbit ethernet was going over the top, but I did have a couple of boxes that had fast ethernet capability, so I am in the middle of upgrading all my systems to run 100Mbit, along with installing a desktop 100Mbit switch. So what if my WAN connection is only 128Kbits, at least between ourselves we can communicate 760 times faster. I note that in the last month or so, our local exchange has now been upgraded to offer ADSL, but I figure that I would have to run as an encrypted VPN back to the cisco office, so ISDN is still an attractive option.

Another recent addition has been a CD burner. As an owner of two digital cameras, I recently was forced to consider that losing over 2 gigabytes of photos would be tantamount to burning all the photo albums I have along with the negatives. Once I had all those photos on a couple of nicely packaged CDs I was able to breathe a sigh of relief.

After a while ignoring the laptop, it got swapped for an even shinier one, an IBM Thinkpad with an obscene amount of memory and disc space. I figured that it would be just criminal to run Windows on it, so after so messing around with partitions (cisco IT people _really_ don't like you to play around with their standard systems), I got a real operating system installed. I had to keep it dual boot, of course, since (like my other home machines, one FreeBSD for real work, and the Windows one for marketing stuff) I figured that no matter what I did, I still needed to run some Windows applications.

Even after installing StarOffice to do the Office style stuff, a lot of our engineering documents are in FrameMaker. So I figured I still need _some_ Windows box lying around. Just as I was about to capitulate to this awful thought, I discovered this magical piece of software called Vmware, which allows a virtual PC to run a guest operating system under a real operating system, Oh joy! So now I have 7 computers, 2 of them virtual.

So, I figure I have it worked out now. I have a laptop running FreeBSD for use at the office and while travelling, which runs Windows 2000 as a guest operating system for when I want to run Framemaker and other applications, and I also have this environment running on my home FreeBSD machine. Each person in my household gets a computer, Sherilyn (spouse) to surf the net planning holidays in expensive island resorts and printing digital photographs of said holidays, son and daughter sending e-mail to friends and surfing, claiming to actually do school projects on their computers (now I know how and why my manager used to have that funny skeptical look on his face). Tying all this together is a 100Mbit switched network, of which I strenuously deny all claims of overkill.

Of course, now that it's all operational, my wife has been looking at Real Estate web sites. Never fear, there is no need to even think about all those dusty ceilings and trying to get Cat5 cable through that wall space. A year or two ago, cisco happened to acquire this really neat wireless networking company called Aironet. No more cables for me! Perhaps I'll get to be a beta user of the new 54Mbit wireless systems. Hmmm, sometimes there _are_ advantages to being employed by the biggest networking company in the world...

Politically Correct UNIX: System VI Release notes

UTILITIES

1) "man" pages are now called "person" pages.

2) Similarly, "hangman" is now the "person_executed_by_an_oppressive_regime."

3) To avoid casting aspersions on our feline friends, the "cat" command is now merely "domestic_quadruped."

4) To date, there has only been a UNIX command for "yes" - reflecting the male belief that women always mean yes, even when they say no. To address this imbalance, System VI adds a "no" command, along with a "-f[orce]" option which will crash the entire system if the "no" is ignored. 5) The bias of the "mail" command is obvious, and it has been replaced by the more neutral "gender" command.

6) The "touch" command has been removed from the standard distribution due to its inappropriate use by high-level managers.

7) "compress" has been replaced by the lightweight "feather" command. Thus, old information (such as that from Dead White European Males) should be archived via "tar" and "feather".

8) The "more" command reflects the materialistic philosophy of the Reagan era. System VI uses the environmentally preferable "less" command.

9) The biodegradable "KleeNeX" displaces the environmentally unfriendly "LaTeX".

SHELL COMMANDS

1) To avoid unpleasant, medieval connotations, the "kill" command has been renamed "euthanise."

2) The "nice" command was historically used by privileged users to give themselves priority over unprivileged ones, by telling them to be "nice". In System VI, the "sue" command is used by unprivileged users to get for themselves the rights enjoyed by privileged ones.

3) "history" has been completely rewritten, and is now called "herstory."

4) "quota" can now specify minimum as well as maximum usage, and will be strictly enforced.

5) The "abort()" function is now called "choice()."

TERMINOLOGY

1) From now on, "rich text" will be more accurately referred to as "exploitative capitalist text".

2) The term "daemons" is a Judeo-Christian pejorative. Such processes will now be known as "spiritual guides."

3) There will no longer be a invidious distinction between "dumb" and "smart" terminals. All terminals are equally valuable.

4) Traditionally, "normal video" (as opposed to "reverse video") was white on black. This implicitly condoned European colonialism, particularly with respect to people of African descent. UNIX System VI now uses "regressive video" to refer to white on black, while "progressive video" can be any color at all over a white background.

5) For far too long, power has been concentrated in the hands of "root" and his "wheel" oligarchy. We have instituted a dictatorship of the users. All system administration functions will be handled by the People's Committee for Democratically Organizing the System (PC-DOS).

6) No longer will it be permissible for files and processes to be "owned" by users. All files and processes will own themselves, and decided how (or whether) to respond to requests from users.

7) The X Window System will henceforth be known as the NC-17 Window System.

8) And finally, UNIX itself will be renamed "PC" - for Procreatively Challenged.

[[The editor of AUUGN has not been able to track down the original source for this crusty old dirge. If you know, auugn@auug.org.au]]

Easy Steps to Samba: Linux Orbit HOWTO

John Gowin <jgowin@linuxorbit.com>

Step 1: Where do I get Samba?

Samba is a great tool for letting Microsoft Windows users share hard drives and printers with Linux users. With the recent release of Samba 2.2, the Samba team has made major improvements to their server software, including support for Windows 2000 and NT 4.0 clients. In addition, Samba 2.2 now can be configured as a Windows NT Primary Domain Controller. In the coming series of tips, we're going to show you how to get started using Samba 2.2.

First, you'll need to know where to get it. Most Linux distributions come with a version of Samba on their distribution CDs. If you can't find a copy of Samba on your distribution of Linux, download it from the Samba web site. http://www.samba.org

Step 2: Installing Samba

You can install Samba one of two ways. Both ways depend on the type of file you download from the Samba web site. If you download a binary package for your distribution, make sure to familiarize yourself with where the various Samba files are installed. For example, if you install the Red Hat RPM for Samba 2.2, your configuration files will be installed in the /etc/samba/ directory.

If you download the Samba source files and compile it on your Linux system, Samba will install in the /usr/local/samba directory. The configuration files will be in the /usr/local/samba/conf/ and the executable files will be in the /usr/local/samba/bin directory

The all-important configuration file for Samba is named smb.conf. If you're not sure where this file is on your system, just use the find command below to find where it is. Make sure that you have superuser privileges when you run it in a terminal window.
find / -name smb.conf -print

For those of you that have locate installed on your system, the command "locate smb.conf" will also work just fine.

Step 3: Getting familiar with smb.conf

Open the smb.conf file in your favorite text editor to make configuration changes. Although there are some graphic interfaces that you can use to configure Samba, you should familiarize yourself with the configuration file smb.conf when configuring Samba for the first time.

You'll notice that there are two types of comment lines in the smb.conf file. Any line that begins with a "#" or a ";" is not recognized by the Samba server daemon. For lines that describe or define a particular setting in the configuration file, you'll see that the line begins with a "#". Lines that begin with ";" are usually valid configuration settings that aren't currently being used.

The smb.conf file is broken down into 2 sections: the Global Settings and the Share Definitions. First, we'll make some changes to the Global Settings.

Step 4: smb.conf Global Settings

At the top of the smb.conf file, you'll find the section labeled Global Settings. First, you'll want to edit these settings at the top of the Global Settings

```
# workgroup = NT-Domain-Name or Workgroup-Name
   workgroup = MYGROUP # server string is the
equivalent of the NT Description field
   server string = Server description
```

The workgroup setting is for configuring the NT Domain or Windows Workgroup that your Samba server will be included in when you browse your Windows Network Neighborhood. If you're configuring Samba for your home network, the workgroup MYGROUP should work just fine.

It should be noted that if you intend to share network devices on your Linux machine from Windows, you'll need to configure your Windows machine for "Client for Microsoft Networks" in your Network Neighborhood properties. This is not the default setting for many consumer Windows systems.

The server string setting will be the description of the server when you browse it in the Windows Network Neighborhood. This can be set to anything you like, but it's useful for beginners to use this setting to identify the server. For example, setting the server string to Samba 2.2 Server will make it easy to find from Windows.

Step 5: More Global settings in smb.conf

The fastest way to configure Samba to share a Linux drive with a Windows user is to create a login on both systems with identical login names. This will become clearer when we talk about Share Definitions. For now, we need to talk a little bit about passwords.

In your text editor, skip down to the setting in your smb.conf file that reads

```
# You may wish to use password encryption. Please
read
# ENCRYPTION.txt, Win95.txt and WinNT.txt in the
Samba documentation.
# Do not enable this option unless you have read
those documents
;  encrypt passwords = yes
;  smb passwd file = /etc/samba/smbpasswd
```

Remove the two semi-colons from the bottom two lines if you're going to connect to Samba from Windows 98/NT/2000/ or ME. If you're connecting from Windows 95, you don't have to edit these two settings.

By default, Samba uses plain text passwords. This was fine when Windows 95 was the predominant end-user operating system for shared Windows machines

because Windows 95 also used plain text passwords. Later versions of Windows use encrypted passwords. Rather than doing some nifty Windows registry editing to work around this problem, turn on Samba's encrypted passwords by uncommenting these two lines.

Once you've removed these semi-colons, you can then skip down to the Share Definitions section of your smb.conf file.

Step 6: Share Definitions

By default, your smb.conf file sets up Linux user home directories to share with Windows machine logins with the same user name. At the top of the Share Definitions section of your smb.conf file, you should see these settings

```
[homes]
   comment = Home Directories
   browseable = no
   writable = yes
```

Change the browseable setting to "yes" to enable browsing of the user's home directory from a Windows machine. Once you've changed this, save your changes to smb.conf and exit your text editor. Now you just have to create a Samba user password for your share.

Step 7: Setting your Samba user password

If you configured Samba for encrypted passwords as we did earlier (Step 5), you'll need to add a user to your smbpasswd file. The username should match your Windows login name and password (case--sensitive). To add a user to the smbpasswd file, you'll need to use the smbpasswd command. If you installed Samba from a binary package, you can simply run the command

```
smbpasswd -a username
```

in a terminal window as the superuser root. If you installed Samba from source, run the same command in your /usr/local/samba/bin/ directory with a "./" in front of the command. you'll be prompted for a password for the user you're creating. Enter the password that you use under Windows for the username you're adding.

That's all you need to do for configuring Samba to share a drive with your Windows network computer. Of course, you'll need to run the Samba server first.

Step 8: Running the Samba server daemon

Using our quick configuration for sharing a Linux drive, all you need to do now is run the Samba server daemon on your Linux machine to enable the shared drive. Once again, your installation choice will determine where you can run the server daemon executable. Binary install users should look for the smb script created in your /etc/rc.d/init.d or /etc/init.d. Simply run the script as the superuser root with the command

```
./smb start
```

If you installed from source, you can run the server daemons smbd and nmbd from the /usr/local/samba/bin directory with the commands

```
./smbd -D
./nmbd -D
```

This starts the Samba server. So why the two commands? Well, if you really must know - read the next step.

Step 9: The server daemons smbd and nmbd

For those of you who just have to know why Samba has two server daemons, the answer is simple. Networks can be complex. Each server daemon allows the use of a different network service for sharing devices over a network. By using both server daemons, you provide services for all types of Windows machines.
The smbd server daemon provides the file and print services to SMB clients, such as Windows 95/98, Windows NT, Windows for Workgroups. SMB stands for "Server Message Block" and is defined as a network protocol for sharing files, printers, serial ports, and communications abstractions such as named pipes and mail slots between computers.

The nmbd server daemon allows for NetBIOS over IP name service requests over a network, like those produced by SMB/CIFS clients such as Windows 95/98/ME, Windows NT and Windows 2000.

If you don't care about this stuff and just want to get your Samba share working, not to worry. We're almost there.

Step 10: Sharing a device from Windows

Once you've run the Samba server daemons on your Linux machine, you should be able to login on your networked Windows machine and see the Linux computer name in your Network Neighborhood. (Remember, you should have the same username and password on both the Linux and Windows machine.) If you do see the Linux computer name there, open Windows Explorer (not Internet Explorer) and open the Tools menu. Select the Map Network Drive option. This will open a dialog box for your share definition.

Choose the drive letter you want the network drive to use (E:, F:, etc.) and then type in the Path to the shared drive. This should look like this

```
\\LINUXCOMPUTERNAME\USERNAME
```

The LINUXCOMPTERNAME should be the name of the Linux computer as it appears on the network and the USERNAME should be your username. It might take a moment or two, but you should then see the new drive letter in Windows Explorer with the home directory of your username on the Linux machine shared.

Step 11: Where to look when things don't work

It never fails. You follow the directions to perfection and the darn thing doesn't work. This tip is for those

who've followed our directions for configuring Samba for a simple share and it didn't work.

First, if you can?t map the network drive with Windows, go to your Linux machine and open a terminal window. Change users to the superuser root with the su command. Remembering how you installed Samba and change directories to the /usr/local/samba/bin/ directory (if you installed from source) or the /usr/bin/ directory. Run the program

```
./testparm
```

This program will read your smb.conf file and tell you if there are any settings configured incorrectly. If you find no errors there, check your Samba log files. On most systems, the Samba log files will be kept in the /var/log/samba/ directory. The log files will be named log.smbd and log.nmbd. Check both of them for any errors.

Last but not least, check your Linux system logs in the file /var/log/messages. You may see smbd or nmbd errors here as well.

Step 12: Using a Windows printer from Linux

Well, now that you've tested your knowledge of Samba, perhaps you're ready to take on a different problem. In the next few steps, we'll show you how to use a Windows shared printer from your Linux machine.

The first thing to do to use a Windows printer is to make sure that the printer is an actual network shared device in Windows. If it's not, you won't be able to use the printer from Linux at all.

On your Windows computer, open your Settings menu from the Start menu. Select Printers to list the printers configured on your system. When you see your printer in the window that opens, right-click the printer's icon and select Sharing from the menu. This will bring up a dialog box with two radio buttons and three data entry fields. Make sure that the Shared radio button is selected and take note of the printer name in the first data entry field. If the printer doesn?t have a name, enter a name in the first field and write it down.

That's basically all the preparation you need on the Windows side. Now that you know that the Windows printer is shared, you're ready to tackle the Linux configuration to use it.

Step 13: Configuring Linux to use a Windows printer

The first step in configuring your Linux system to use a Windows shared printer is to create an entry in your /etc/printcap file. If you've never used a printer on your Linux system, this file should be empty.
There are two ways to go about adding a printer device to your Linux system. You can edit the file manually and add a printer to your configuration or you can use a program that helps you configure it. If you've never done this before, we recommend you use a utility program to help you. The syntax for the /etc/printcap file can be a little tricky if you've never seen it before. In this case, we?ll use the Red Hat printtool utility. you'll need to run printtool in a terminal window with X running.

When you open the printtool utility, select the Add button to add your Windows printer. From the menu that appears, select the radio button that says "SMB/Windows 95/NT Printer" and click OK. Now you should see the Printer Entry dialog box. The first three fields in the form will be filled out for you with these settings

```
Names: lp
Spool Directory: /var/spool/lpd/lp
File Limit in Kb (0 = no limit): 0
```

You can't leave these as they are. you'll need to fill in at least three on the fields that are left. First, enter the name of the Windows machine in the Hostname of Printer Server field. If you use IP addresses on your network, you can enter the IP address of you Windows machine in the IP number of Server field. This field is optional. Only use it if you know the IP number. Enter the name of the printer as it was listed in the Share menu on the Windows machine in the Printer Name field. If the printer name was in all capital letters, be sure that you enter it here the same way.

The last thing you should do is select a printer filter.

Step 14: Selecting a Printer Filter

When we left you in our last step, we were ready to select a Printer Filter in the Red Hat printtool utility. Click the Select button next the heading Input Filter. This will open a menu that will list the available printer filters. If you're not sure what type of printer yours is, check the make and model and see if it is in the menu list. If not, select a printer filter that is similar to your printer, especially if the make is the same, but the model is not. Once you choose your filter, don't worry about the other settings you see on this menu. You can tweak these later when you know that the printer filter you've chosen works.

Once you've chosen a printer filter, click OK until you return to the printtool main menu. Close the printtool menu and take a look at your /etc/printcap file. You should see settings that look like this:

```
##PRINTTOOL3## SMB printerfiltername
lp:\
:sd=/var/spool/lpd/lp:\
:mx#0:\
:sh:\
:if=/var/spool/lpd/lp/filter:\
:af=/var/spool/lpd/lp/acct:\
:lp=/dev/null:
```

The printerfiltername will correspond to the one you chose in the printtool utility. As you can see, without doing a lot of reading on how to create /etc/printcap entries manually, you?d find it difficult to get it right the first time without using printtool.

Step 15: Testing the printer

To test the shared printer from Linux, try printing any text file with the lpr command in a terminal window. Try the command

```
lpr /etc/printcap
```

to print your printcap file. The file should print, but if it doesn't, check the following:

- Is your Windows machine on, the printer online and have paper?
- Are the Samba server daemons running on your Linux machine?Run the Samba utility testprns to make sure all entries in your /etc/printcap file are correct. Use this especially if you manually edit the file.
- Check your Samba logs for errors.

If printing to the Windows printer does work, you should be able to print from all of your programs. Open a browser and try to print a web page for the final test. If the web page prints, you're ready to go!

Step 16: What shares are available on your network?

End users always seem to get left out of the information loop. If not left out, they're the last to know. If you know that Samba is configured on a network you use, there's a simple Linux command that you can use to determine what shares are available from a Samba server.

Of course, Samba has to be installed on your Linux system too, but you'd have to have it installed to use a shared Samba network device anyway.

If you know the name of a Samba server on your network and want to know what devices it's sharing, run the smbclient utility. It should be located in either the /usr/local/samba/bin or /usr/bin directory on your system. Run this command in a terminal window to find out what a particular Samba host is sharing with the network.

```
smbclient -L sambahostname
```

Step 17: Getting help with Samba online

Don't feel too bad if you don't get your Samba configuration right the first time. Many experienced Linux users have trouble, and it's mainly because there are so many places that you can go wrong. Using the correct syntax in your configuration files is the usual culprit, so always double-check your work with the testparm and testprns utilities.

You can also find a lot of Samba information online. The documentation for Samba is thorough, but can be over the head of beginners at times. Always check your man pages for Samba as well. you'll find all variables for your Samba configuration here along with their definitions.

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

http://www.linuxorbit.com/howto/sambahowto.php3

KOffice 1.1 Beta 3 Review: Part 1

Author: Kent Nguyen - kent@mslinux.com
Editor: Tina Gasperson tina@gasperson.com

For five years, the only office suite I knew and used was Microsoft Office. Logically, an upgrade to Office XP should be my next step. Yet, after reading about Microsoft's increased invasion of privacy with XP (Microsoft to me: We're turning off your Office), it became my quest to find an alternative office suite.

I stumbled across KOffice when a Linux guru friend suggested it to me. Personally, I wouldn't want to switch applications, let alone an operating system, just to use an office suite. Yet now I was stuck between upgrading to Office XP or making the bold move to install Linux and Koffice 1.1 Beta 3. I thought, "why not?" since Linux will always be free, no one person has control over it, and I won't be harassed five years later. And so my quest began.

My quest:

I called upon a friend to help me setup Mandrake Linux 8.0 and KOffice 1.1 Beta 3 on my system. It took him about 30 minutes to have everything up and running. When I booted up, the first screen looked like the logon screen I see in Windows NT at work. At home, I have Windows '98, which doesn't require ahave a logon screen. Obviously, I need to login to Linux. It's kinda like Windows NT in that way.

My first impression was that in KDE, everything looked very much like Windows. I was relieved. The "K" is identical in function to the "Start" button in Windows, so that was where I clicked first. I scanned the list of software categories to find "KOffice." After several minutes scratching my head, wondering why I can't find KOffice, I asked my guru friend. He showed me that KOffice was under "Office". Ah! After moving my mouse over "Office", the menu expanded to "KOffice Workspace", "KWord", "KSpread", "Kivio", "KChart, "KIllustrator", "Krayon", and Kugar". All these names were new to me, but after several seconds of thought, I made some good progress associating those names with what I used to use. For instance, KWord compares to Microsoft Word, KSpread to Excel, Kivio to Visio, KIllustrator to Adobe Illustrator, and Krayon to Adobe Photoshop. So much for the guessing game. It was time for me to get down to business.

KWord:

I clicked on "K"->"Office"->"KWord" to launch KWord. Up came a pretty screen with four options to choose from. I chose to create a new document, and clicked OK.

KWord loaded quickly. The first screen looked like this:

The first document I created with KWord was a short memo to management about this exciting new software called KOffice.

The memo demonstrated basic word processing features such as center, sizing text, changing fonts, tabs, and word wrap. One feature I stumbled across that I like is the font dialog drop down. It allows me to preview the font.

Microsoft Word has a feature that I find very helpful: it keeps the four most recently used fonts at the top of the dropdown. I couldn't find this option in KWord; it's probably a setting I need to enable. After writing the memo, I wanted to see what it looked like before I printed it KWord has a couple of ways to do this. One is the really good "Preview Mode". You simply go to "View"->"Preview Mode".

In addition to the preview mode, there's a "Print Preview". The print preview allows you to see what it will look like if it is printed. You can go to "File"->"Print Preview" to see a preview of your document.

Another awesome feature most users from the Microsoft world would enjoy is the ability to view "Formatting Characters". You go to "View"->" Formatting Characters".

Though I haven't gone through the advanced features of KWord, like frames, I'm impressed by what I've seen so far. Great job KWord team!

I'd been meaning to write a letter to Charles Schwab for quite some time. Mr Schwab owns a large stake of the investment company that's named after him. As a part of this review, I composed a letter telling him how wonderful KOffice is, and asking him to pass the word to Bob Lee, the infrastructure desktop manager. I clicked on "File"->"New" and selected "US Letter" template. The template set everything up for me. It created the layout as 8.5" by 11" and defined the global measurement as inches instead of millimeters. I created a frame by clicking the "ab" icon on the frame toolbar. The red line points to the "ab" icon.

When prompted to give a name to the frame, I just clicked "OK" instead of entering a name. Once that was done, I moved the cursor over the frame, where it changes from a pointer to a compass-like icon. Then I pressed "Ctrl-C" to copy, and "Ctrl-V" to paste, and used my mouse to drag and resize the dimensions of the frames according to my needs.

After I put all the frames in place, I started to write my letter by clicking inside the frame. Here is the finished letter with visible frames in place.

I discovered that KWord has this neat feature that allows me to make the visible frame border disappear, which made my document look very nice.

To do this, I clicked on "View"->"Frame Borders". The resulting document looked like this:

After I finished spell checking the document by clicking on "Tools"->"Spelling...", I was ready to print.

When I first heard of frame-based word processing, I was a bit afraid to use it. Yet, KWord has made frame-based word processing a reality for the end--user. KWord has done it this time! From what I have seen so far, it surpasses what Microsoft Word can do.

In the next part of this review, I will delve deeper into the power of KWord as a desktop publishing tool, an area in which Microsoft Word lacks. This is where I think KWord really stands out against the competition.

While visiting the offices of our local newspaper, The Stockton Record, as a child, I saw a few Unix machines and one Mac. I'm not sure what software they used on the Unix machines, but I think it was Framemaker or something.

On the Mac it was Adobe Pagemaker. I thought it would be cool to have the capabilities they have.

I didn't have to go far to know that KWord has these capabilities. In newspaper publishing they have what is called a layout, which is simply the way the dimension of the paper is set. I set my layout at 14" in width and 23" in height by creating a new document and using the option "Start with an empty document".

The default page dimensions and units are 8.5" x 11" and in millimeters. To change the default page unit, I put the cursor over the ruler bar, right clicked and chose "inches." To change the dimension of the page layout, I went to "Format"->"Page .." and chose "Custom". From there I could enter the specific width and height I wanted to use - 14" x 23." Page...

This was where I found that frame based word--processing makes desktop publishing a real cinch. I was able to create all sort of frames in my new 14" x 23" layout.

To preview it without the frame border, click "View"->"Frame Border".

As you can see, the power of a frame-based word processor allowed me the flexibility to create desktop publishing quality work. This really gives KWord an advantage over Microsoft Word.

The only two competitors in the market that can measure up to KWord frame-based technology are Framemaker and Pagemaker, both of which are Adobe products. Adobe wants to be the de facto of desktop publishing. In my opinion, if KWord makes inroads in the desktop publishing market, Adobe may want to acquire KWord - but as we all know, KWord is GPL and cannot be bought. Because of this, I think within the next 5 years Adobe will lose most its desktop publishing market share to KWord.

KWord is not only frame-based but component-based. The component technology that ties all KOffice applications together is called KParts. Within a KWord application, you can have a KSpread, a Kivio, and a KPresenter document. You can scale, zoom, and do pratically anything to these embedded KParts.

I called up five days ago to request a $100,000 loan to start my new business helping people migrate out of Microsoft and into Linux. The loan officer told me at the very least I have to have a summary report of my current financial circumstances. I took this opportunity to explore the power of component technology in KWord. In other words, I wanted to embed a spreadsheet into my KWord document.

KWord is one of the few mature word processors that has embedded technology. And, it is the only word processor that leverages the power of KParts.I started out writing my financial summary report like so:

The next thing to do was to embed a KSpread into KWord. All I had to do was click on the "?" icon represented by the translucent red line in the enlarged picture above. I clicked on to create an embedded spreadsheet. I noticed a crosshair mouse cursor, which allowed me to specify the side of my embedded KSpread. I was then prompted to either use an existing KSpread or create a new KSpread. I chose to create a new KSpread.

I now had a KSpread spreadsheet in my KWord. I moved the mouse over the spreadsheet, and wow! My KWord application turned into a KSpread application. That, everyone, is the power of KParts -- the most advanced component technology on this planet.(Yup, better than DCOM, I will tell you why later.)

After I finished entering my bogus financial information, I moved the mouse cursor out of the KSpread frame, and clicked on the KWord document. Then I created the border by selecting the KSpread frame. This is different from selecting the KSpread object and editing it. The way KWord works is to place the KSpread object in a frame, which gave me the flexibility to resize, add borders, and layer other frames on top of it. This is what makes KWord state-of-the--art. This technology will uniquely position KWord to be the champion amongst all other word processors.

After working so hard on my KSpread, I wanted to see how it looked in print preview.

As mentioned before, KWord is a frame-based, component-based word processor.

This combination of features makes KWord the most advanced of its kind to date. KWord is chock-full of features, and not only does it use the most advanced technology available, the price AND freedom cannot be beaten. KWord is FREE and GPL.

I guess my short letter to Mr Charles Schwab seemed insignificant when compared to my excitement about how great I think KOffice is. So I decided to share my excitement in a bigger way. I was going to write to some of the most prominent leaders of the world, like President George Bush of the United States, President Johannes Rau of Germany, President Tarja Halonen from Finland, Prime Minister Jens Stoltenberg from Norway, President Wim Kok from Holland, Prime Minister Tony Blair of the United Kingdom, Prime Minister Poul Nyrup Rasmussen of Denmark, President Vladimir V. Putin of the Russian Federation, King Albert II and Queen Paola of Belgium, President Jian Zeming of

China, Prime Minister Yoshiro Mori of Japan, Prime Minister Goran Persson of Sweden, President Jacques Chirac of France, President Vicente Fox of Mexico, Secretary-General Kofi Annan of the United Nations, and 499 CEO of the Fortune 500 companies. This would be the most important letter in my life. I had to make it perfect. There was no room for mistakes. These leaders could easily spot the slightest grammar, spelling, and structural mistake.

That was why I needed to call upon two of my trusted friends, Le Shang Lin and Peter Suchland for help. The network transparency and document view model of KWord enabled me to make this collaborative effort work. Software with these capabilities is most often referred to as groupware.

I assigned color codes to Le Shang and Peter so that I could keep track of who recommended what changes. I assigned Le Shang the color green, Peter the color red, and myself the color blue. I started by using the same letter template I made for Charles Schwab's letter.

I gave Le Shang and Peter read and write access to my file, and then I saved it on the remote server. This is another area where KWord is different from all other word processors. It's the first and only word processor that has built-in network transparency; that is, I'm able to seamlessly save the file to a remote server. Notice the "ftp://www.mslinux.com/reviews/" to the left of the tool buttons in the "Save As..." file dialog. "ftp://" signifies the protocol. "www.mslinux.com" indicates where the remote server is.

All the KOffice applications make use of network transparency, which is made possible by a technology called IO Slave. The KDE team is the pioneer in Internet desktop innovation. They have consistently wowed the general public with technologies that make life a lot easier, a lot happier, and a lot smoother. We often don't see what they've done, but we do enjoy the fruits of their labor. I got everything setup and I was ready to write my important letter.

After I finished my letter, I could use jabber or email to notify Le Shang and Peter that it was ready to proofread. Le Shang was the first to let me know he was currently proofreading my letter. He simply used his KWord to remotely open my "letterToWorld" document. He made a few comments.

Next Peter decided to take a jab at this and also made a few comments.

I read those comments and finished up my letter. I was ready for it to be sent to every leader in the world. KWord is really exciting software. The groupware features are somewhat limited, but are adequate to do the things I need in a collaborative environment. The network transparency made it possible for my friends who are miles away to help me with my correspondence, making comments so I could make it better. This is only the tip of the iceberg for more exciting new groupware features that will be added to KWord.

KWord is now my primary word processor. Over the past days, I've migrated most of my Word documents over. It's a fantastic feeling knowing that I'm free at last. Although KWord is only at Beta3, it sure feels like a final release. What's Next? There are five more applications in the KOffice family that I will review. Below are previews of each. I will review KSpread next.

KSpread
KIllustrator
KPresenter
Kivio
Krayon

Stay tuned ...

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

http://static.kdenews.org/mirrors/www.mslinux.com/reviews/koffice1.html

I clicked on: "K"->"Office"-->"KWord" to launch KWord. Up came a pretty screen with four options to choose from. I chose to create a new document, and clicked OK.

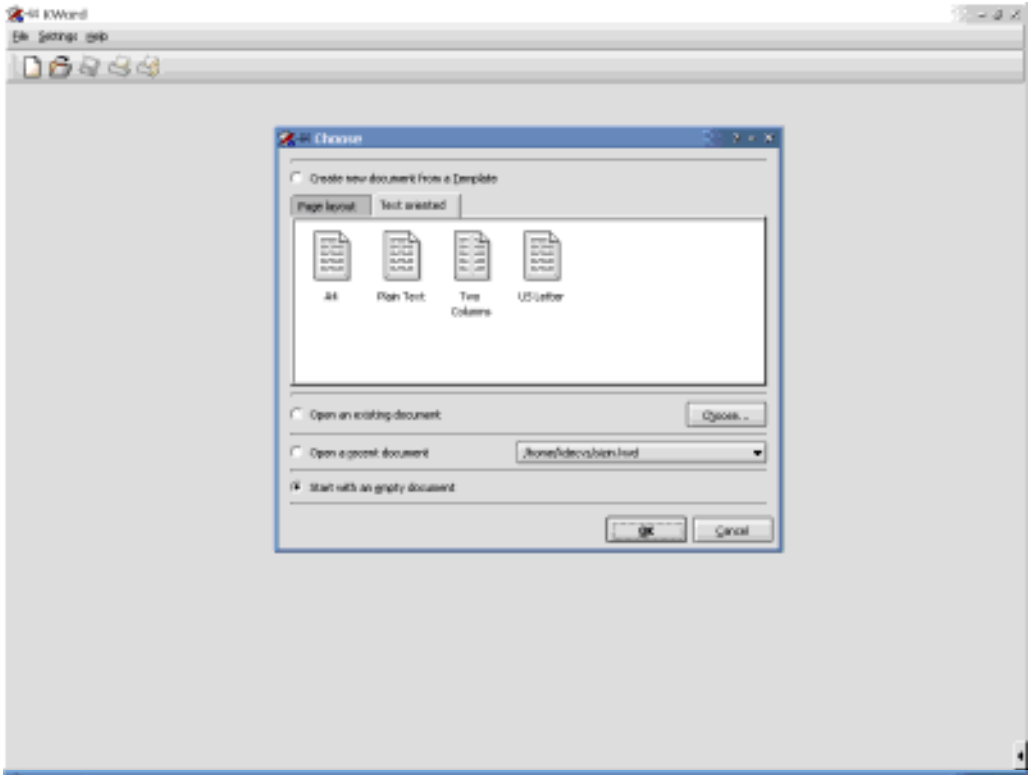

KWord loaded quickly. The first screen looked like this:

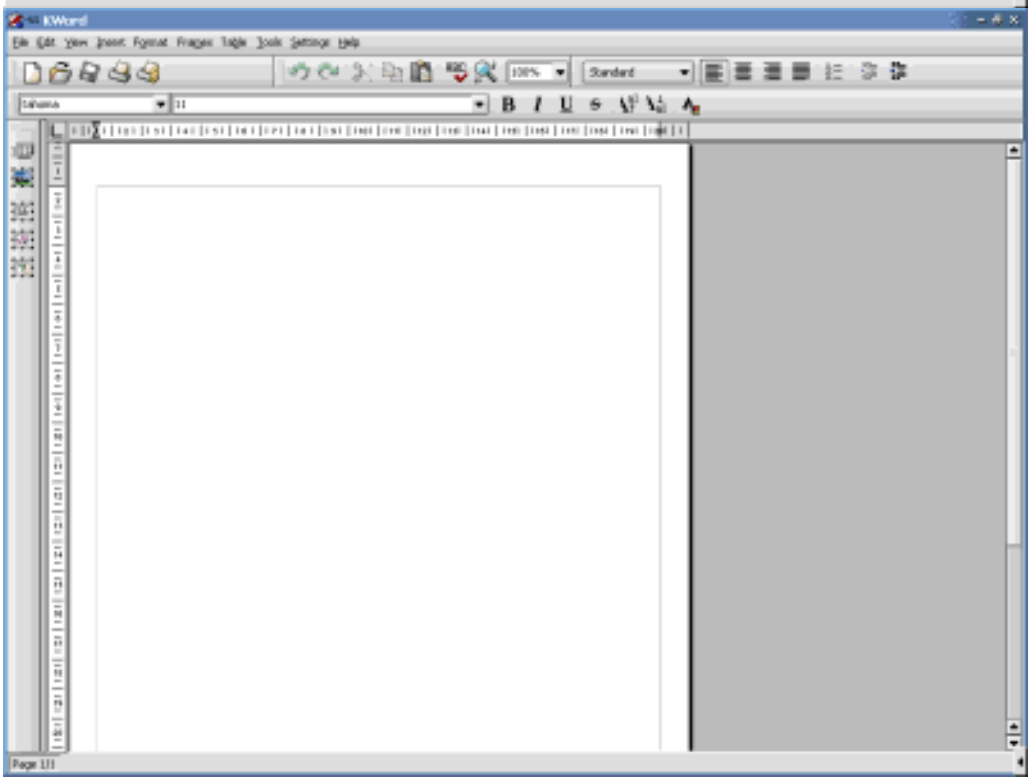

The first document I created with KWord was a short memo to management about this exciting new software called KOffice.

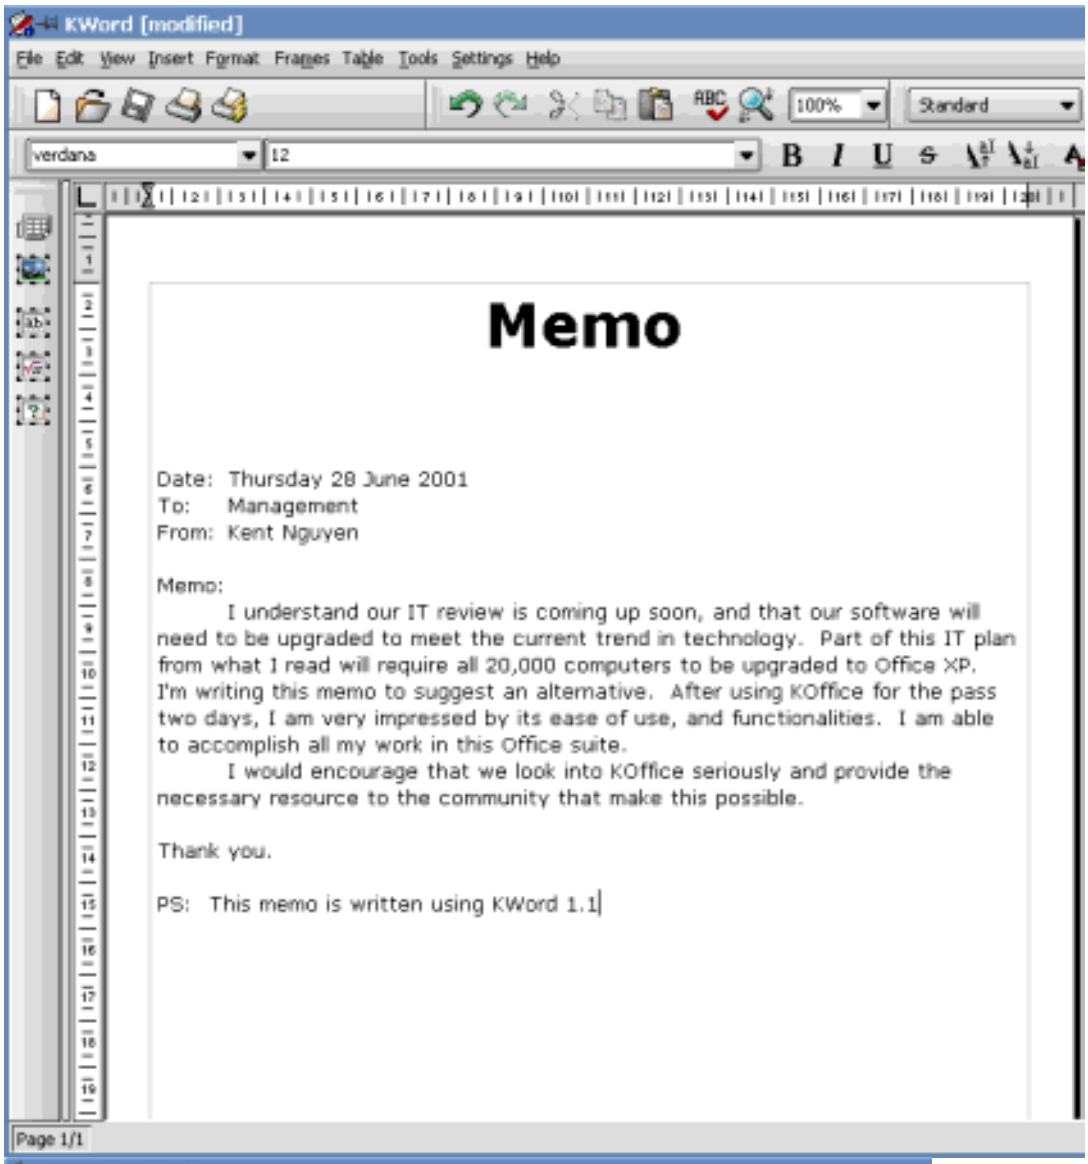

After writing the memo, I wanted to see what it looked like before I printed it KWord has a couple of ways to do this. One is the really good "Preview Mode". You simply go to "View"->"Preview Mode".

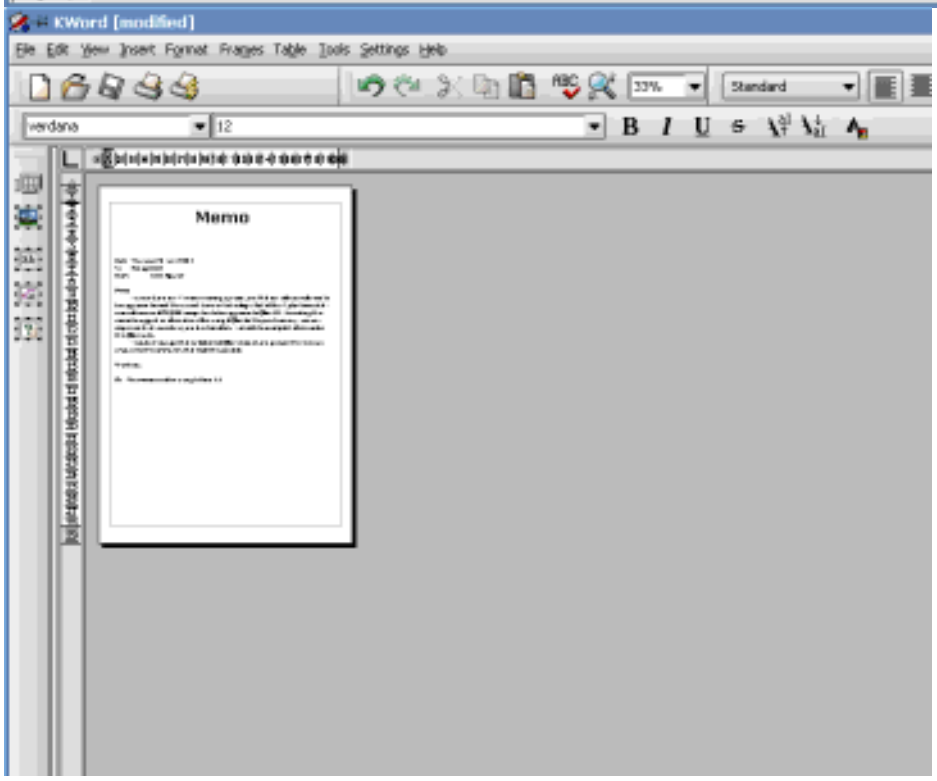

What the final result looks like in a PDF viewer.

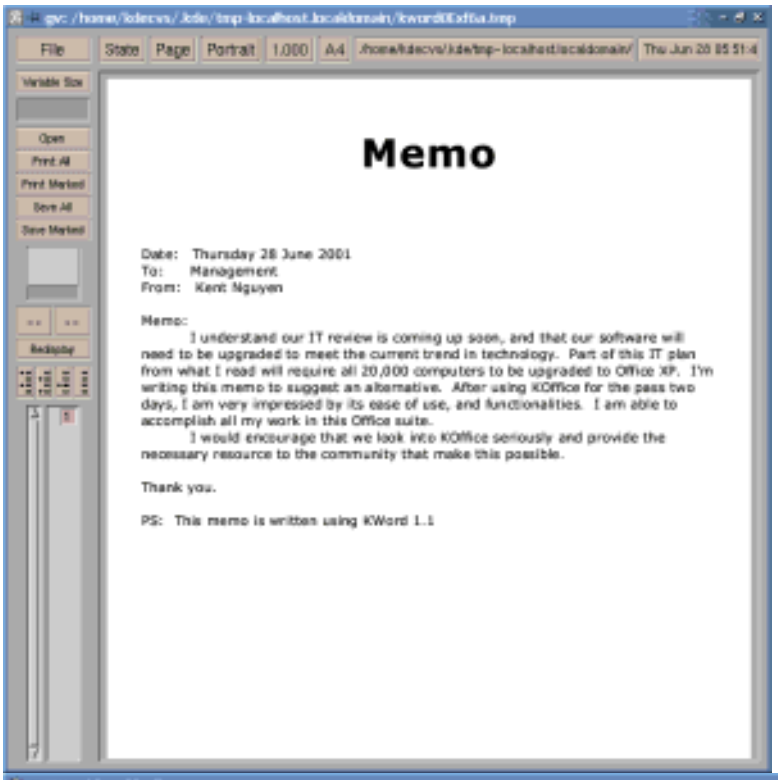

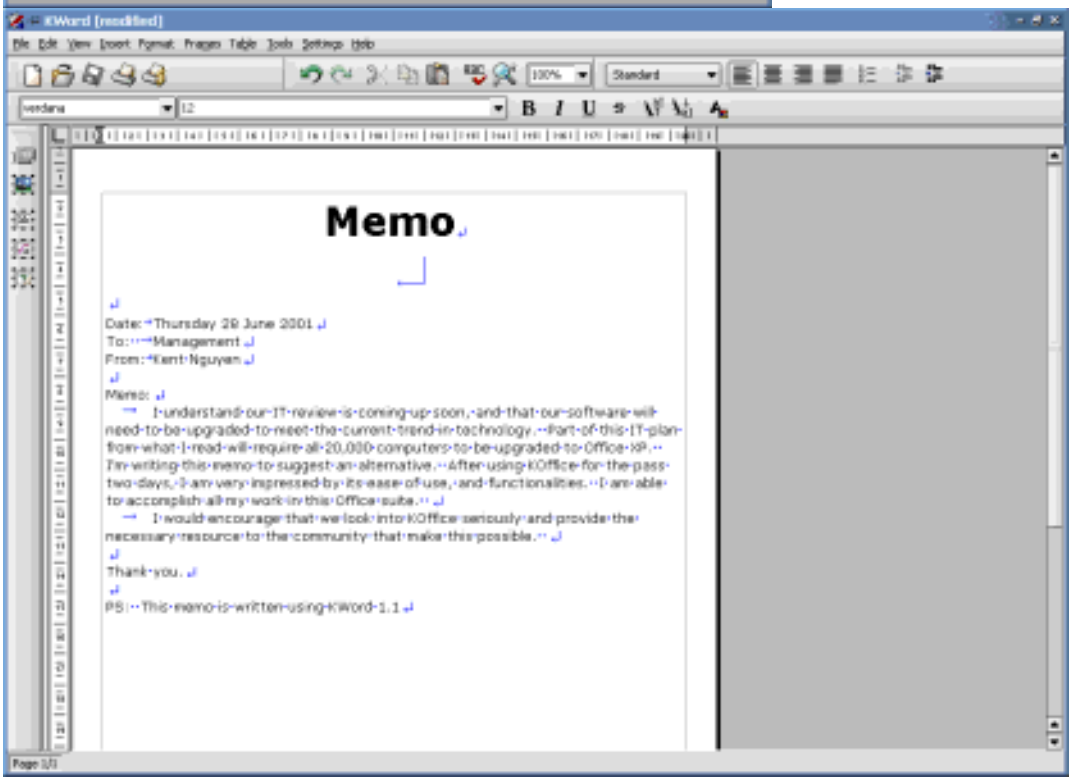

What follows are full-screen snapshots of the various KOffice applications.

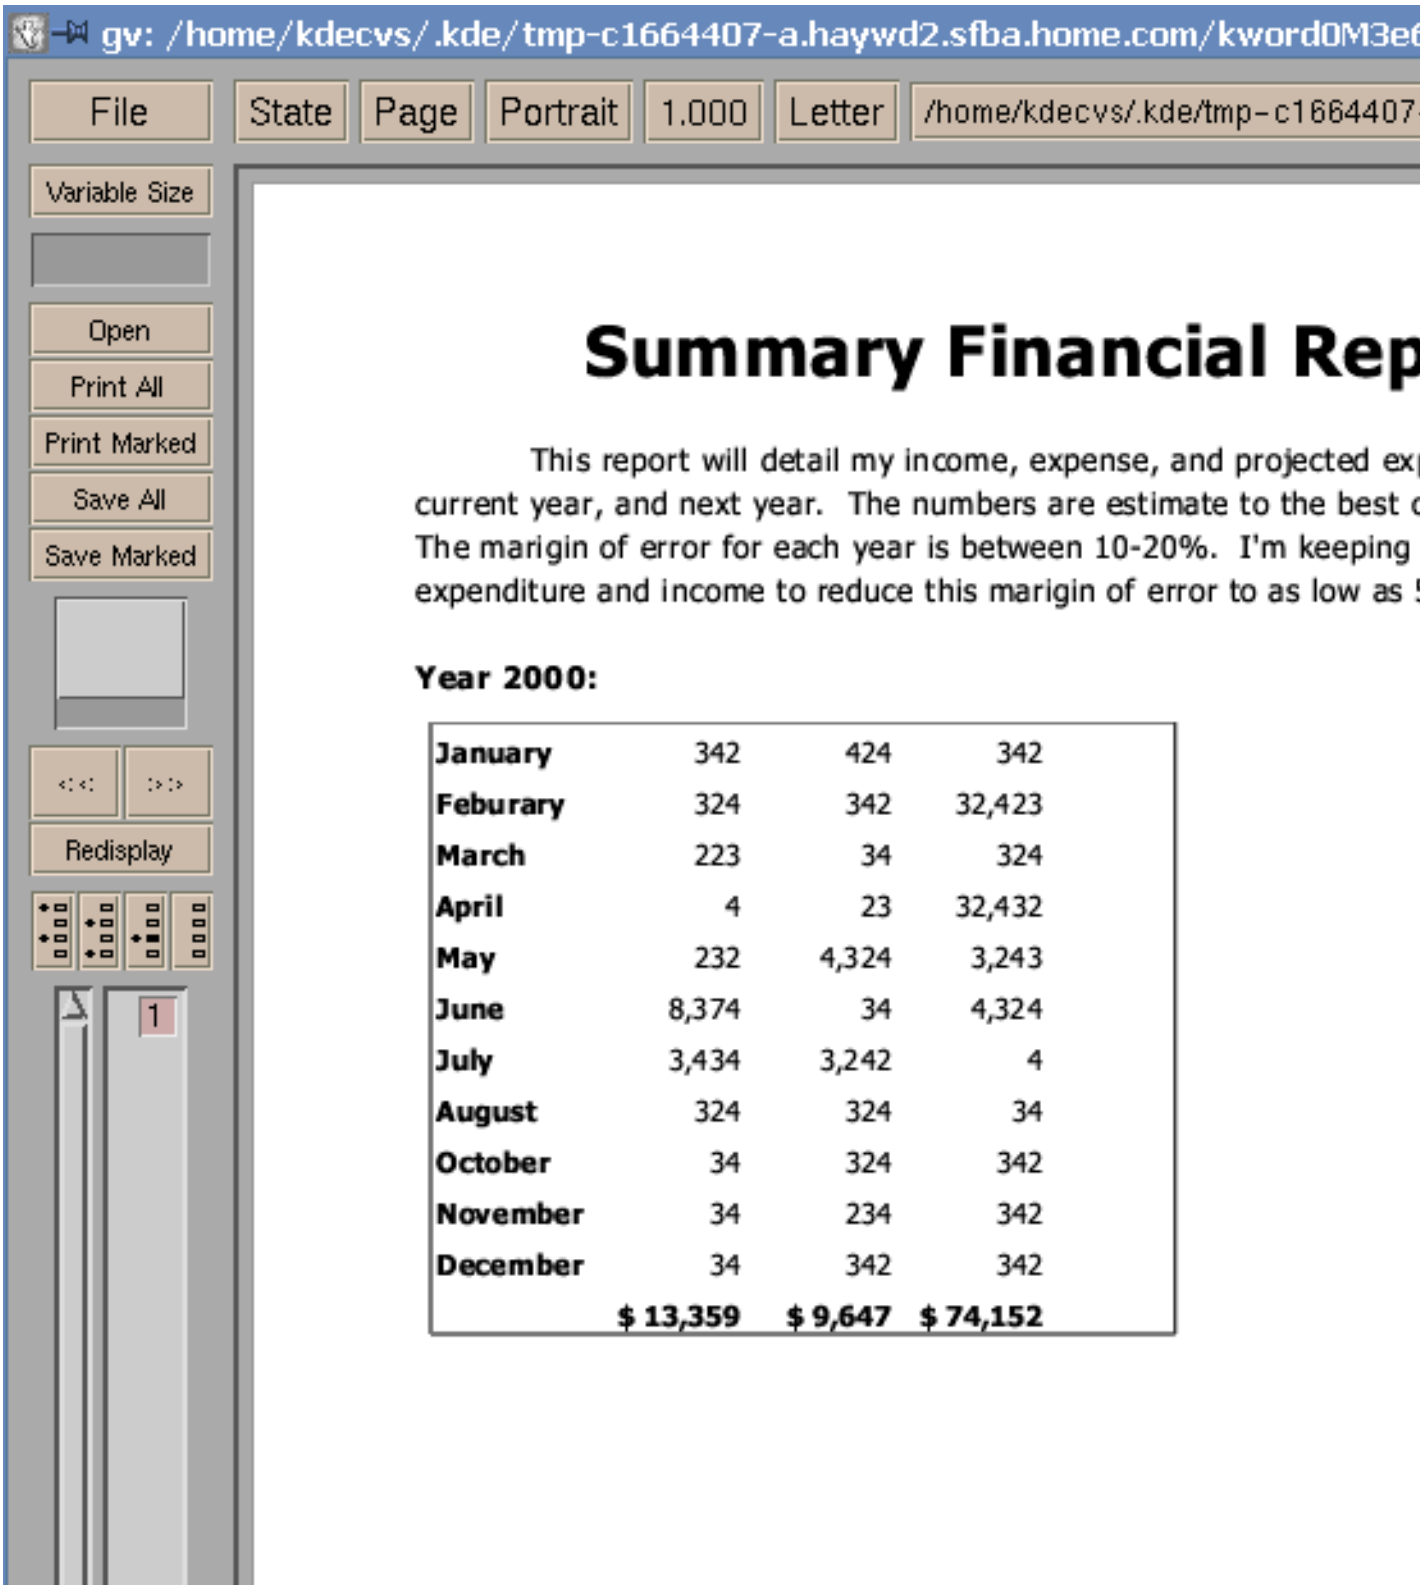
gv: /home/kdecvs/.kde/tmp-c1664407-a.haywd2.sfba.home.com/kword0M3e
File
State
Page
Portrait
1.000
Letter
/home/kdecvs/.kde/tmp-c1664407
Variable Size
Open
Print All
Print Marked
Save All
Save Marked
Redisplay
1
Summary Financial Rep
This report will detail my income, expense, and projected ex
current year, and next year. The numbers are estimate to the best
The marigin of error for each year is between 10-20%. I'm keeping
expenditure and income to reduce this marigin of error to as low as
Year 2000:
January 342 424 342
Feburary 324 342 32,423
March 223 34 324
April 4 23 32,432
May 232 4,324 3,243
June 8,374 34 4,324
July 3,434 3,242 4
August 324 324 34
October 34 324 342
November 34 234 342
December 34 342 342
$ 13,359 $ 9,647 $ 74,152

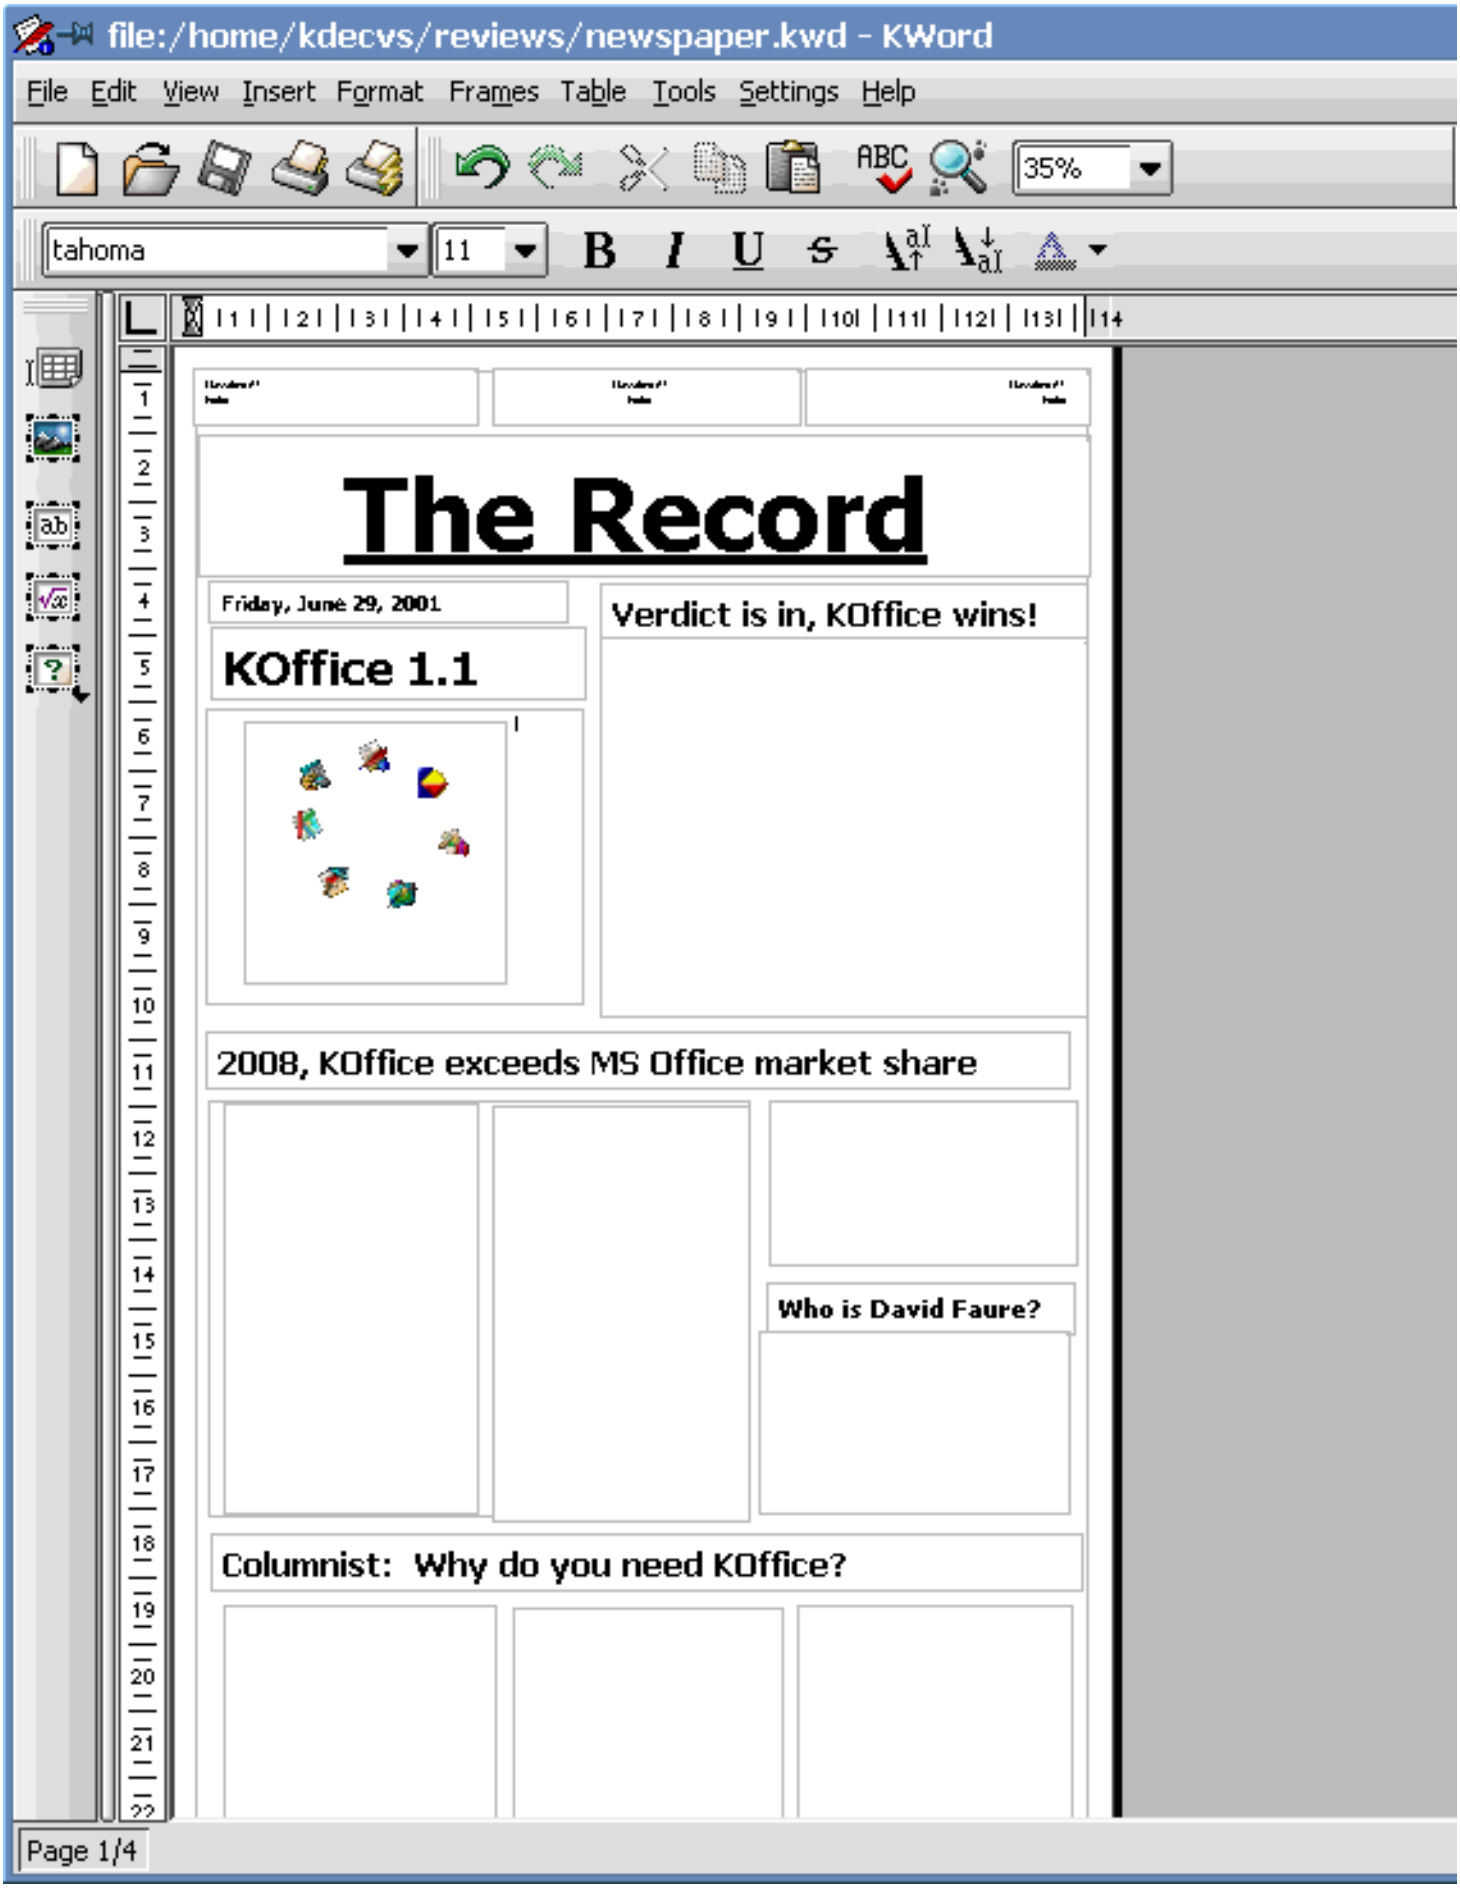
file:/home/kdecvs/reviews/newspaper.kwd - KWord
File Edit View Insert Format Frames Table Tools Settings Help
35%
tahoma
11
The Record
Friday, June 29, 2001
Verdict is in, KOffice wins!
KOffice 1.1
2008, KOffice exceeds MS Office market share
Who is David Faure?
Columnist: Why do you need KOffice?
Page 1/4

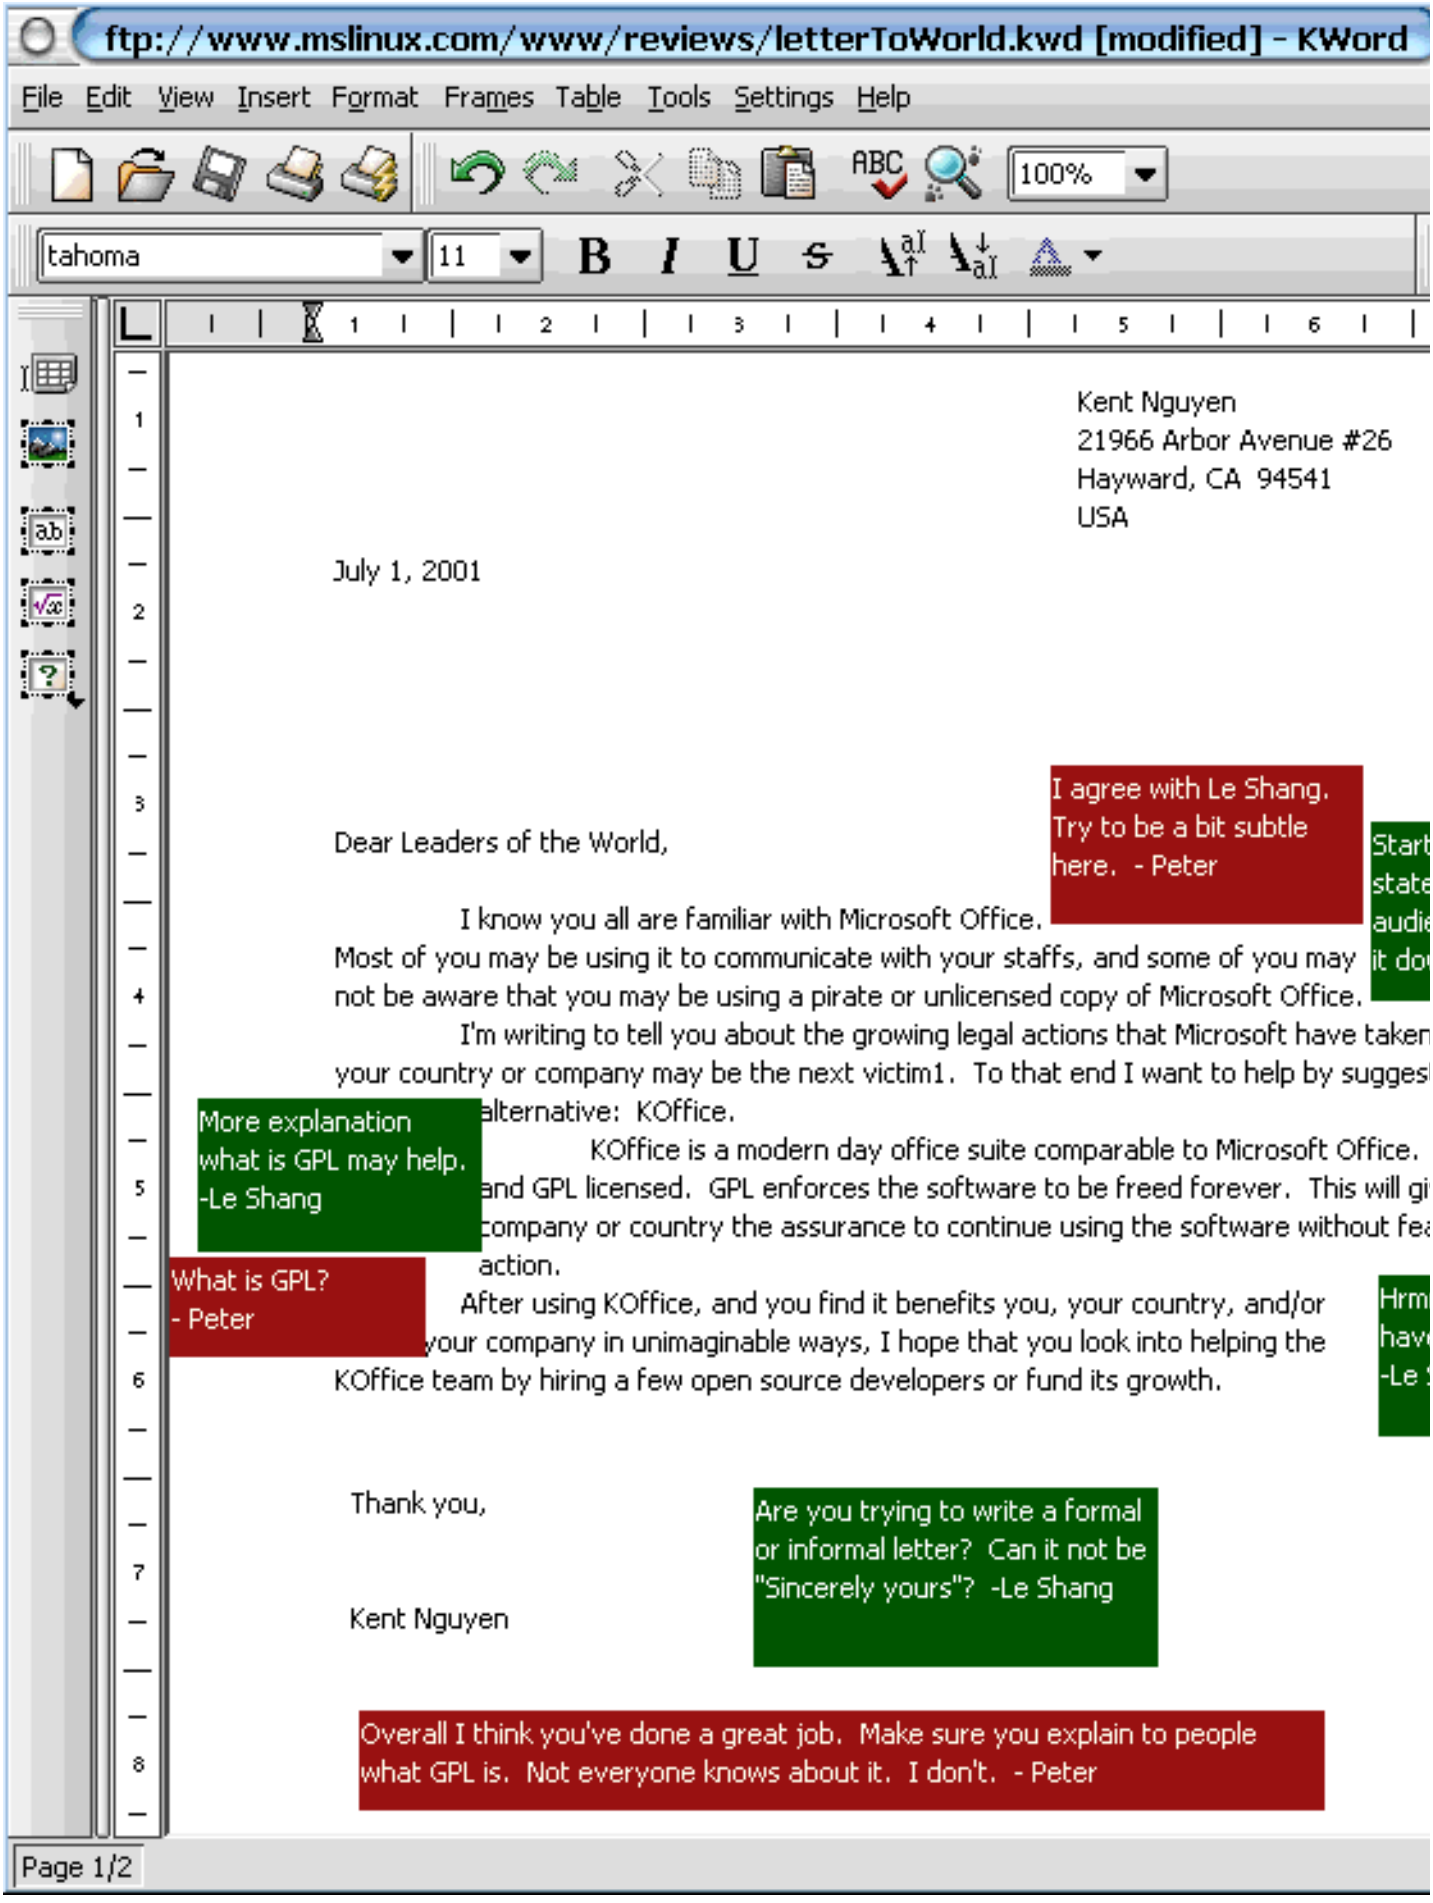
ftp://www.mslinux.com/www/reviews/letterToWorld.kwd [modified] - KWord
File Edit View Insert Format Frames Table Tools Settings Help
100%
tahoma
11
Kent Nguyen
21966 Arbor Avenue #26
Hayward, CA 94541
USA
July 1, 2001
Dear Leaders of the World,
I know you all are familiar with Microsoft Office.
Most of you may be using it to communicate with your staffs, and some of you may not be aware that you may be using a pirate or unlicensed copy of Microsoft Office.
I'm writing to tell you about the growing legal actions that Microsoft have taken
your country or company may be the next victim1. To that end I want to help by sugges
alternative: KOffice.
KOffice is a modern day office suite comparable to Microsoft Office.
and GPL licensed. GPL enforces the software to be freed forever. This will gi
company or country the assurance to continue using the software without fea
action.
After using KOffice, and you find it benefits you, your country, and/or
your company in unimaginable ways, I hope that you look into helping the
KOffice team by hiring a few open source developers or fund its growth.
Thank you,
Kent Nguyen
I agree with Le Shang. Try to be a bit subtle here. - Peter
Start
state
audie
it do
More explanation what is GPL may help. -Le Shang
What is GPL? - Peter
Hrm
have
-Le
Are you trying to write a formal or informal letter? Can it not be "Sincerely yours"? -Le Shang
Overall I think you've done a great job. Make sure you explain to people what GPL is. Not everyone knows about it. I don't. - Peter
Page 1/2

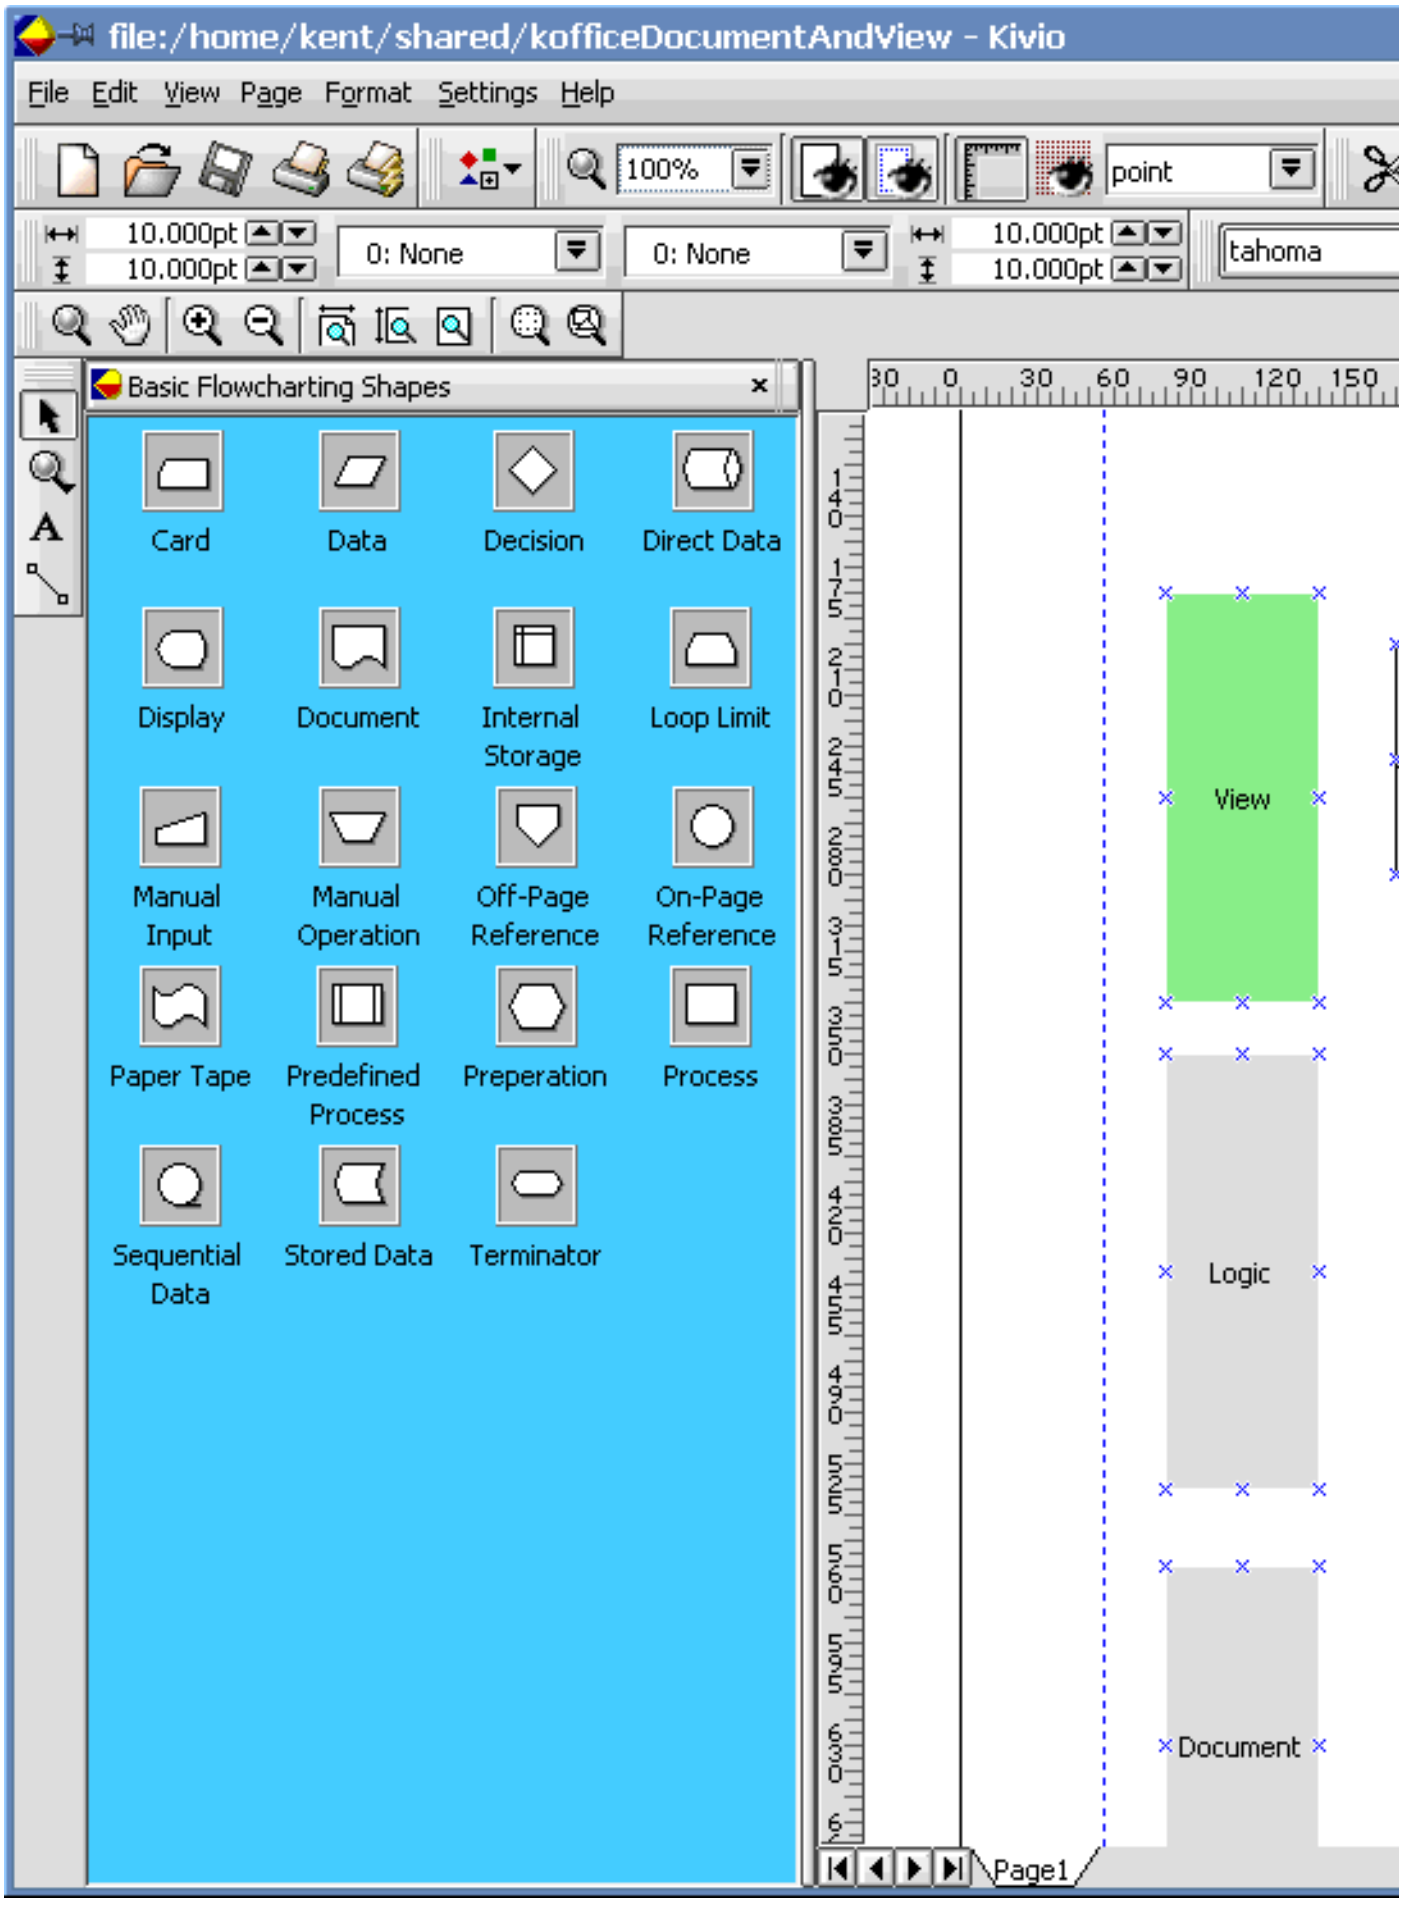
file:/home/kent/shared/kofficeDocumentAndView - Kivio
File Edit View Page Format Settings Help
100%
point
10.000pt
10.000pt
0: None
0: None
10.000pt
10.000pt
tahoma
Basic Flowcharting Shapes
Card
Data
Decision
Direct Data
Display
Document
Internal Storage
Loop Limit
Manual Input
Manual Operation
Off-Page Reference
On-Page Reference
Paper Tape
Predefined Process
Preperation
Process
Sequential Data
Stored Data
Terminator
View
Logic
Document
Page1

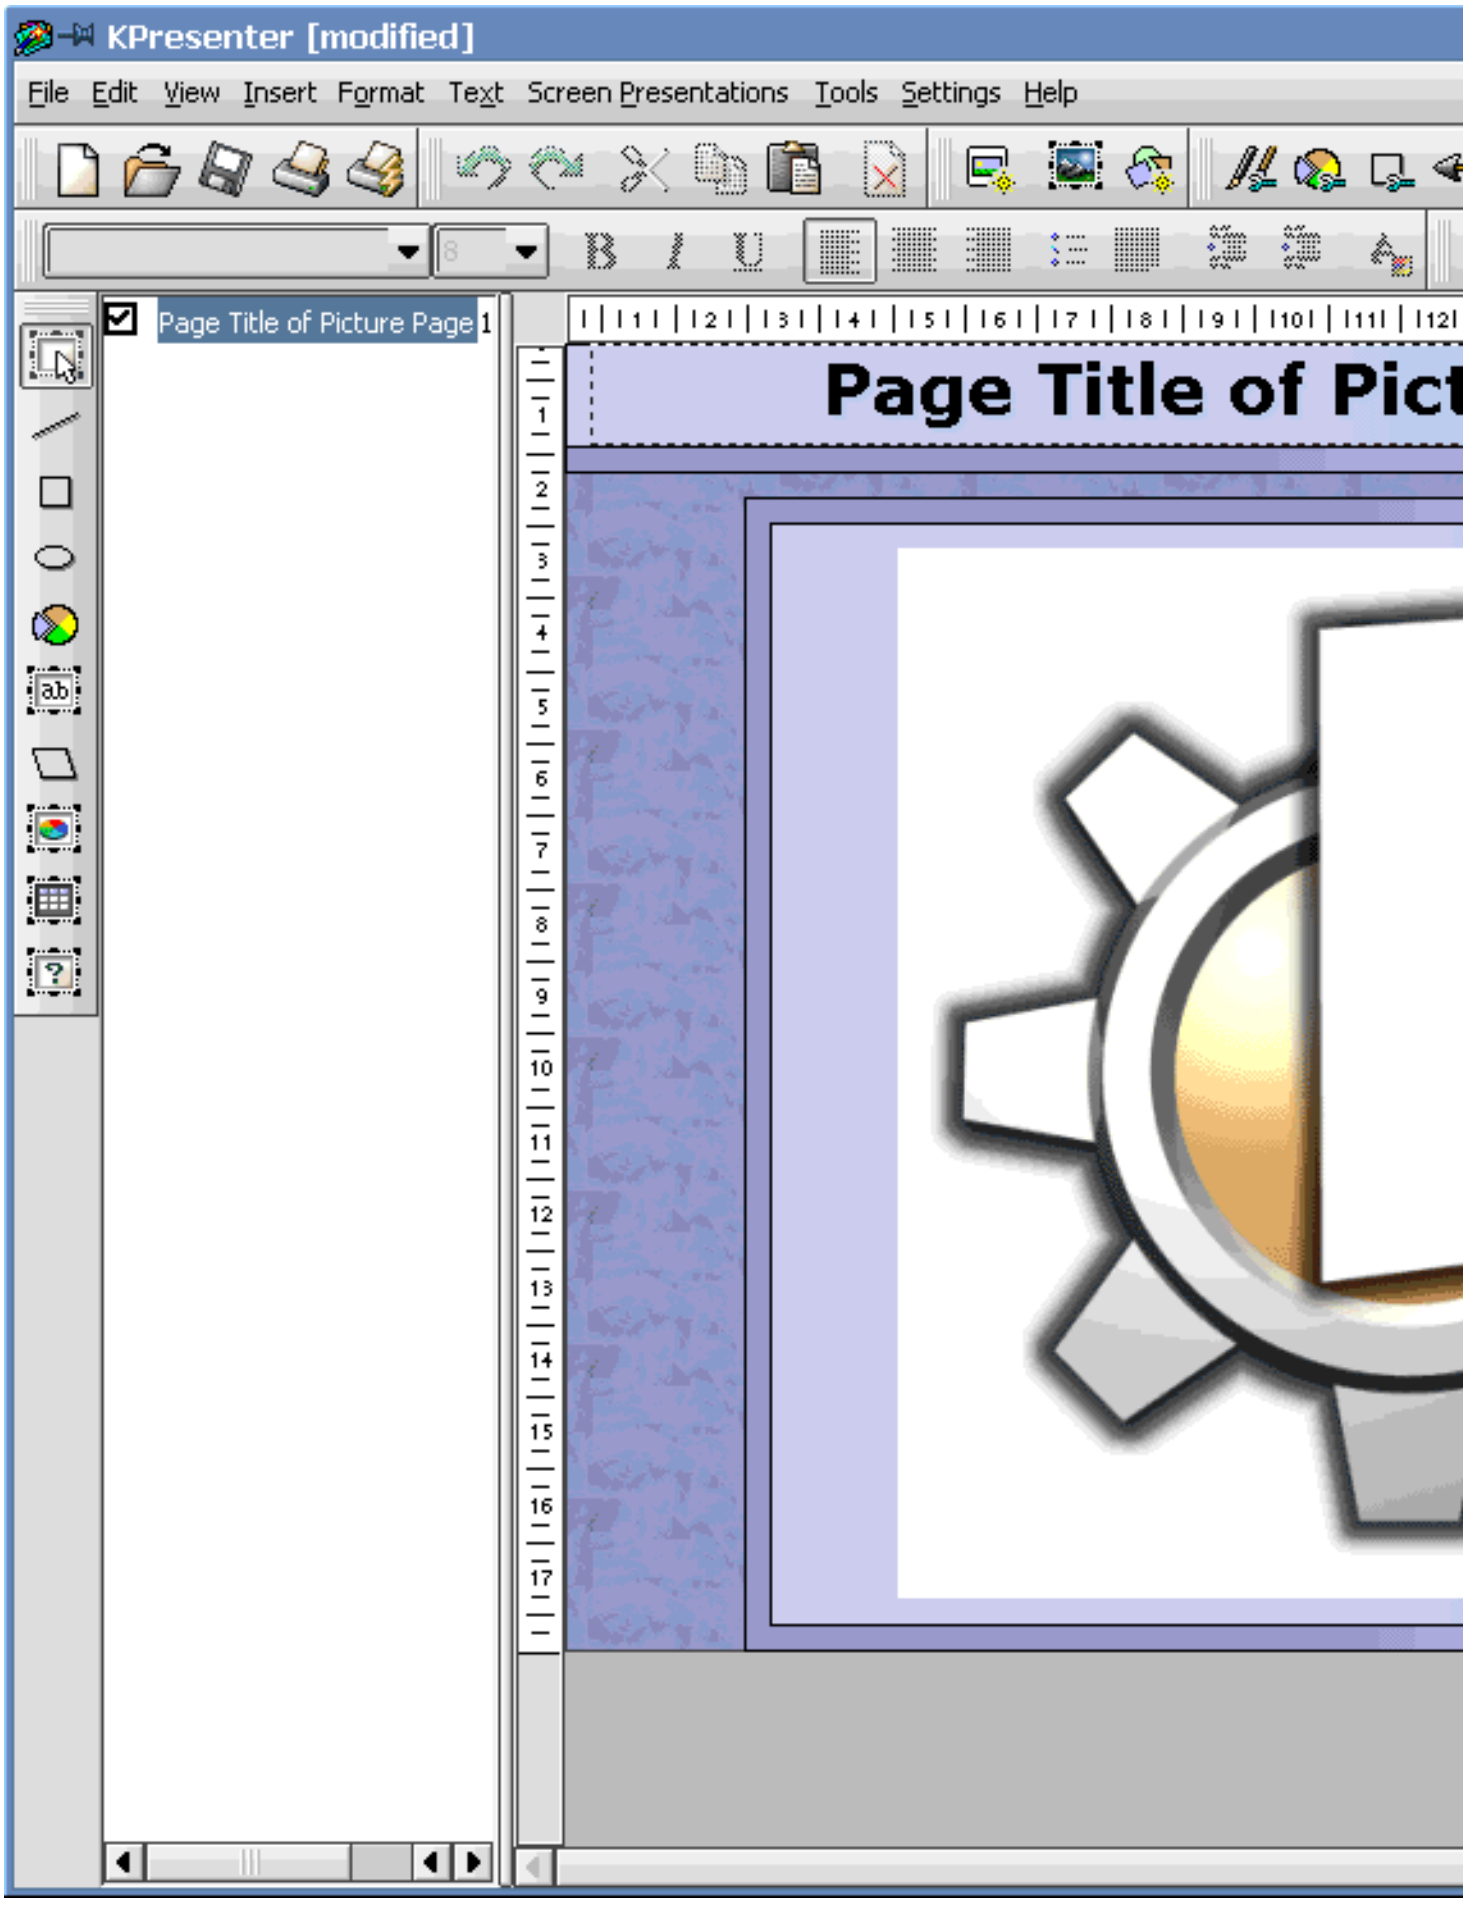
KPresenter [modified]
File Edit View Insert Format Text Screen Presentations Tools Settings Help
Page Title of Picture Page 1
Page Title of Pict

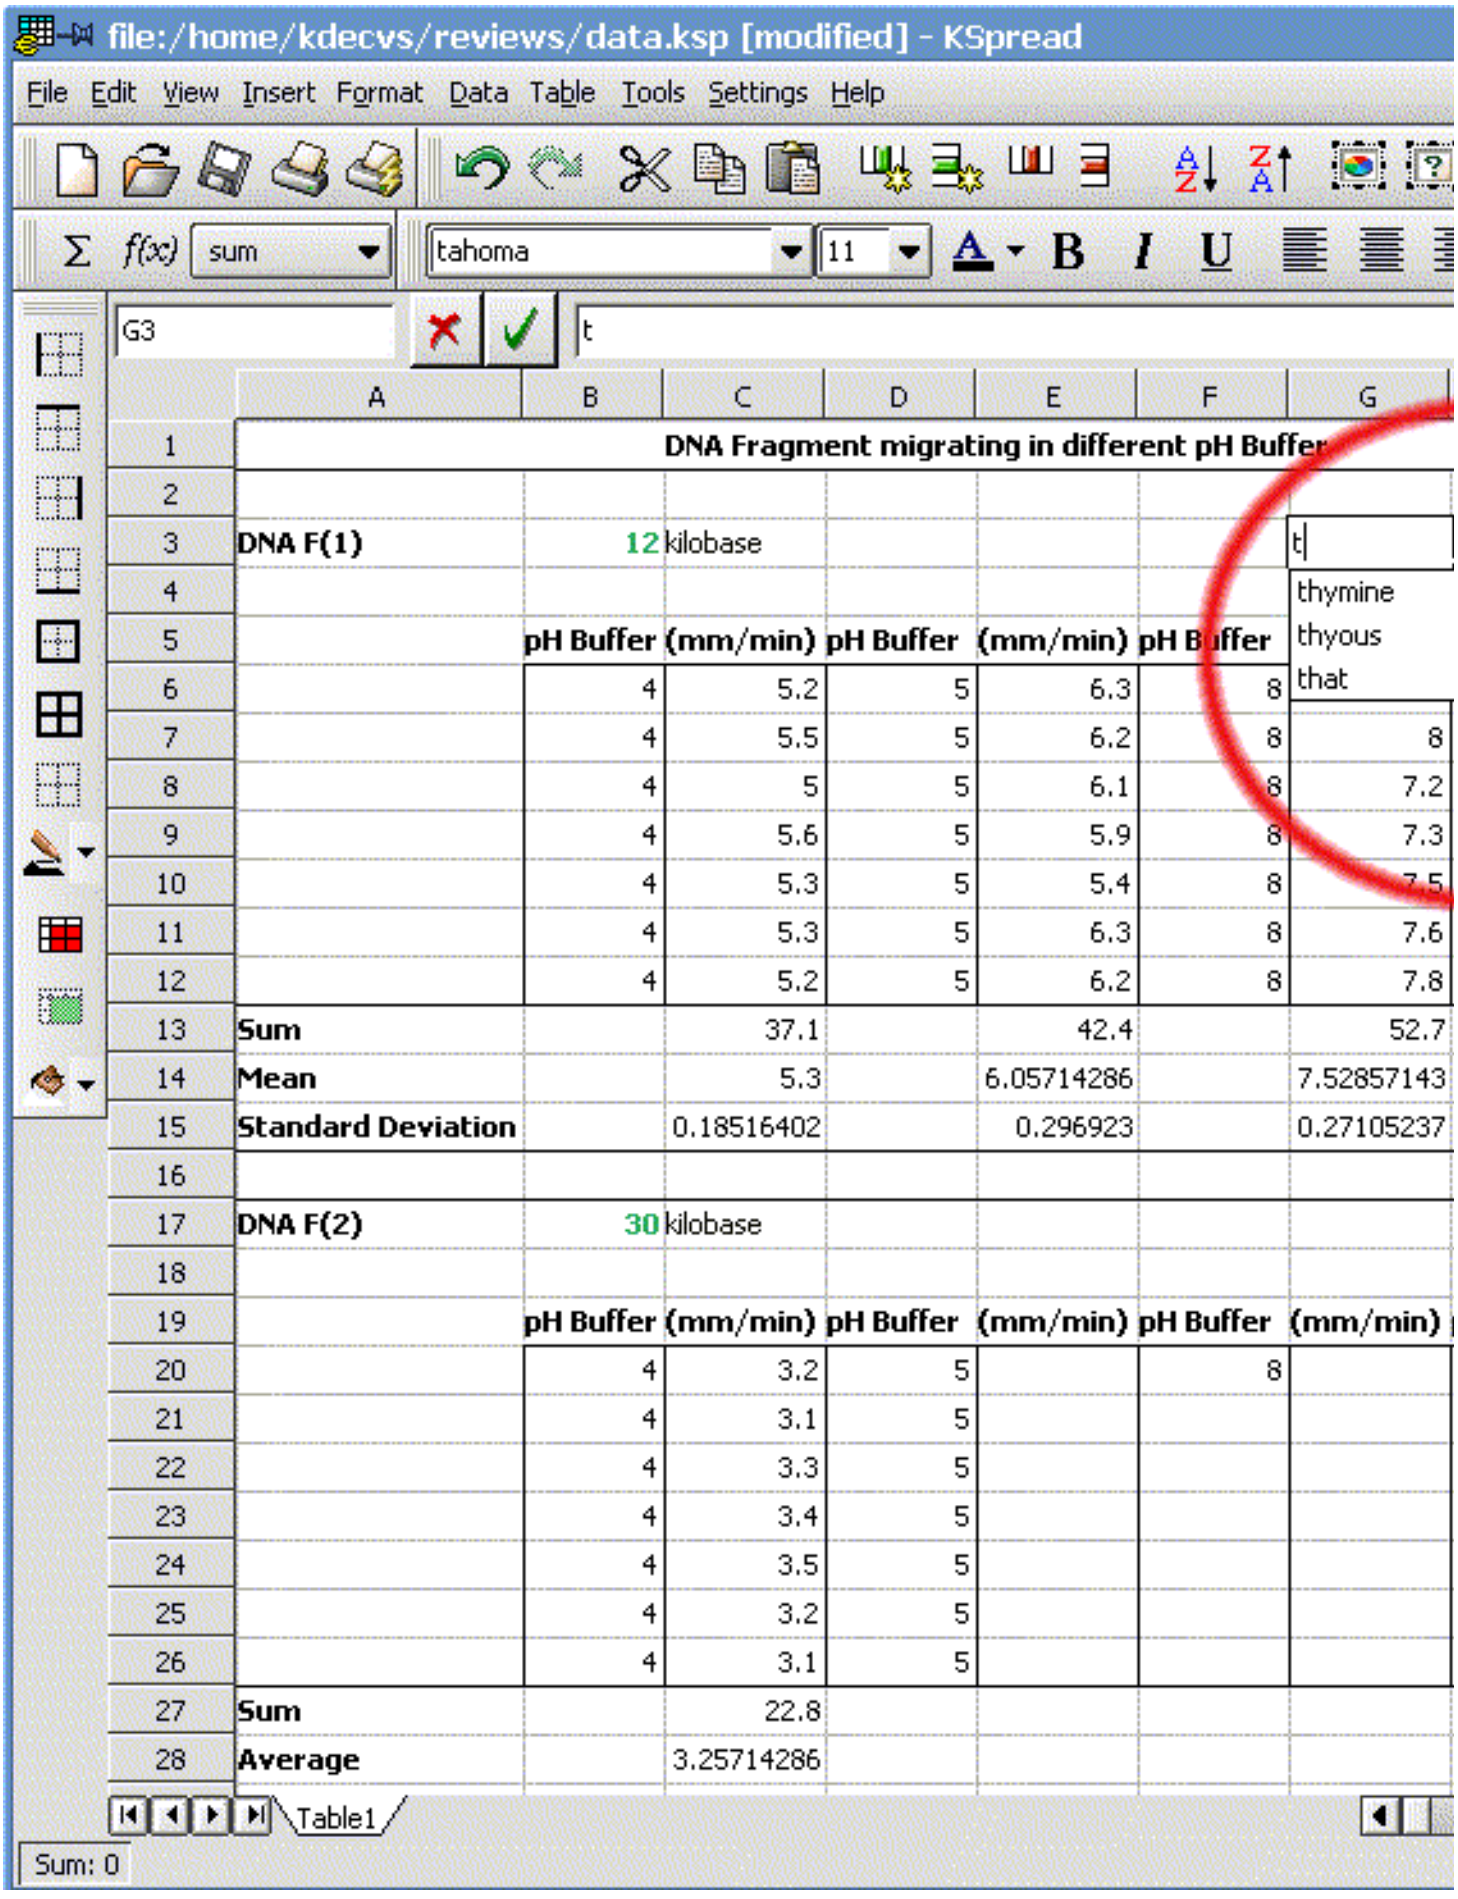
file:/home/kdecvs/reviews/data.ksp [modified] - KSpread
File Edit View Insert Format Data Table Tools Settings Help
sum
tahoma
11
G3
t
DNA Fragment migrating in different pH Buffer
DNA F(1) 12 kilobase
pH Buffer (mm/min) pH Buffer (mm/min) pH Buffer
4 5.2 5 6.3 8
4 5.5 5 6.2 8 8
4 5 5 6.1 8 7.2
4 5.6 5 5.9 8 7.3
4 5.3 5 5.4 8
4 5.3 5 6.3 8 7.6
4 5.2 5 6.2 8 7.8
Sum 37.1 42.4 52.7
Mean 5.3 6.05714286 7.52857143
Standard Deviation 0.18516402 0.296923 0.27105237
DNA F(2) 30 kilobase
pH Buffer (mm/min) pH Buffer (mm/min) pH Buffer (mm/min)
4 3.2 5 8
4 3.1 5
4 3.3 5
4 3.4 5
4 3.5 5
4 3.2 5
4 3.1 5
Sum 22.8
Average 3.25714286
t
thymine
thyous
that
Table1
Sum: 0

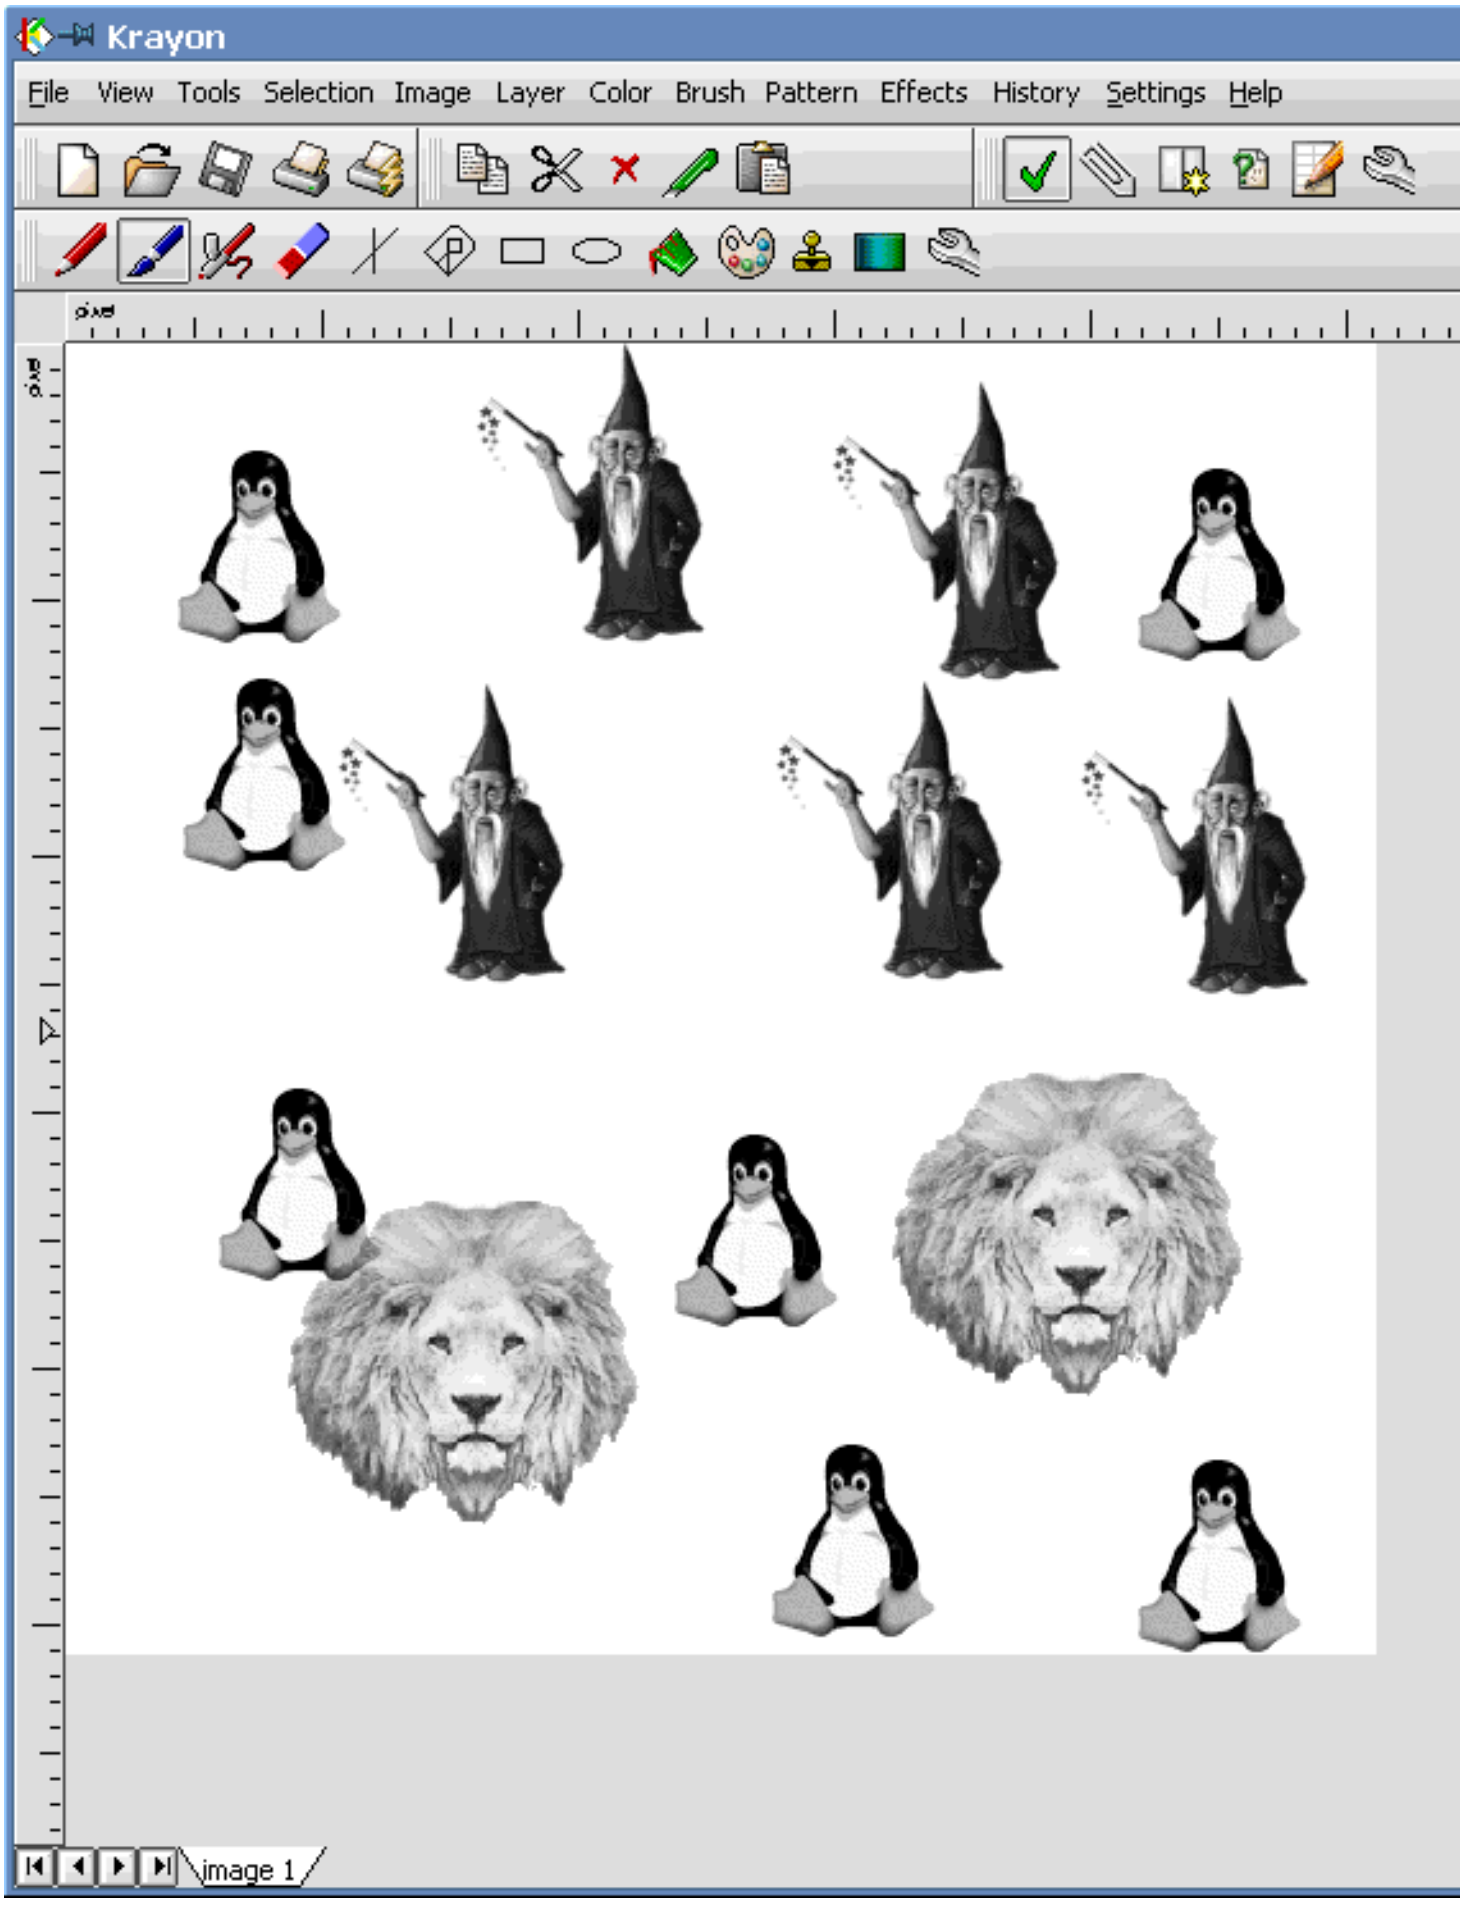
Krayon
File View Tools Selection Image Layer Color Brush Pattern Effects History Settings Help
image 1

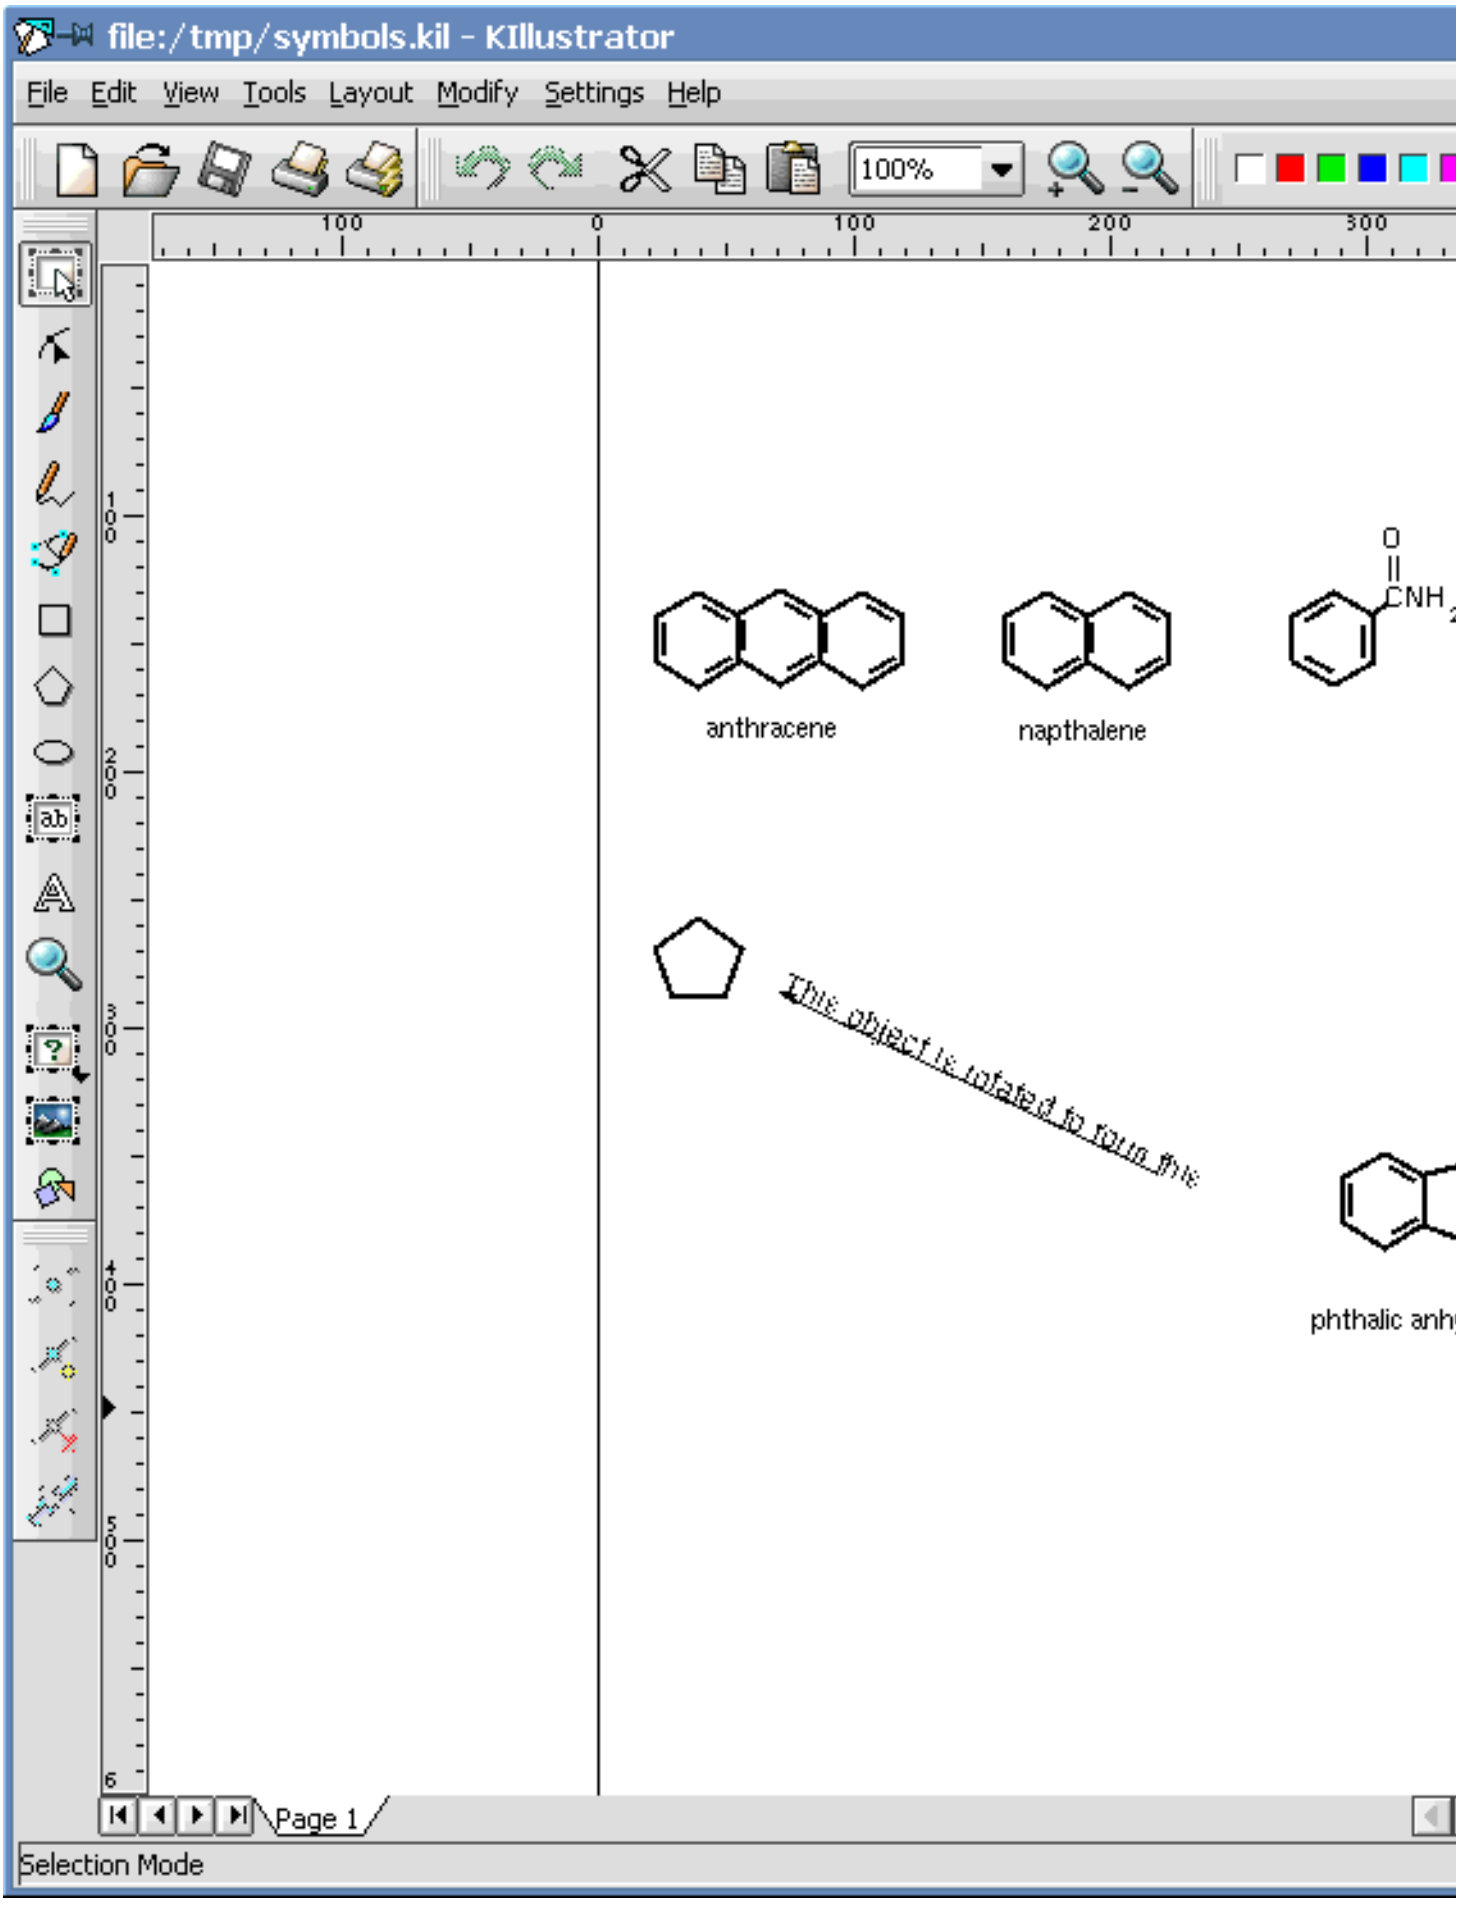
file:/tmp/symbols.kil - KIllustrator
File Edit View Tools Layout Modify Settings Help
100%
anthracene
napthalene
This object is rotated to your fine
Page 1
Selection Mode

Linux-Mandrake 7.2 PowerPack Deluxe Review

Frederik Dannemare <frederik@dannemare.dk>

This review is based on Linux-Mandrake 7.2 PowerPack Deluxe, which was released by MandrakeSoft November 9, 2000.

WHAT DO YOU GET?

With the PowerPack Deluxe edition of Mandrake 7.2 you get all the software you could ever wish for. 7 CDs loaded with more than 2300 apps will probably satisfy most people, I'm sure.

With the PowerPack Deluxe you also get 100 days free e-mail support, a user and installation guide (266 pages), a reference manual (268 pages) and a nice box, of course. Oh yeah, and don't forget the Linux--Mandrake stickers to smack all over your beloved Linux box.

The user and installation guide is well-written and covers topics such as hard disk partitioning, installation with DrakX (Mandrake's graphical installer), using KDE 2, setting up an Internet connection (PPP, ISDN, xDSL and Cable), and hardware configuration with HardDrake.

You have to love all the Mandrake specific "Drak(e) tools" (well, at least I do) such as HardDrake, XFDrake, Userdrake. They make the penguin life a lot easier in many situations for newbies, and advanced users as well.

The comprehensive reference manual takes you through basic Linux/Unix concepts, including console commands, file systems, kernel configuration and compilation.

BASIC SYSTEM SPECS
Kernel 2.2.17 (and 2.4beta for testning), XFree86 3.3.6 & XFree86 4.0.1, glibc 2.1.3, KDE 2 (although, early boxes was shipped with latest KDE2 pre-release), Qt 2.2.1, GNOME 1.2.1, GTK+ 1.2.8, Perl 5.600, Python 1.5.2, etc.
INSTALL MACHINE
Motherboard: Asus P2B
CPU: Pentium II 450 MHz
RAM: 192 MB RAM
Harddisk: Maxtor, 20.1 GB
Graphics: nVidia TNT
Sound: SoundBlaster Live!
Monitor: Samsung SyncMaster 1100p+
DVD-ROM: Pioneer 103S
NIC: D-LINK DFE-530TX
Printer: Canon BJC-2000

INSTALLATION

See the install steps in the image just below.

With the install CD in my drive I boot up my system. Soon after DrakX starts. DrakX is an excellent graphical installer, which I got to know when I tried Mandrake 7.1 for the first time back in June 2000.

But what is now this? My Microsoft PS/2 mouse is not detected by DrakX. Dead mouse is all I get.

This was a disappointment for me, of course, as I have never had mouse problems before with other distributions (except for Corel Linux 1.2, actually).

The DrakX that comes with Mandrake 7.1 also runs like a beauty on my system. This was very odd to say the least.

It is not really an option (at least not to my knowledge) to fully manoeuvre the graphical partitioning tool (which is a part of DrakX) without a mouse, so I have to get it working, unless I want to use text install mode instead. And I don't. I like DrakX better than the text installer included with Mandrake.

Oh well, DrakX probably don't like my MS mouse (I don't blame it), I said to myself. Instead, I tried with a Logitech PS/2 mouse, but with no luck. A third (noname) mouse also wouldn't do the trick, so I mailed Mandrake Support for help, as I was clueless as to what could be causing my mouse problem.

Support mailed me back, but the suggestions I was given didn't bring my mouse back from the dead. Another reply from support also did not solve my mouse problem. I was almost giving up on my review, but then a few days later my phone rang, and Mandrake support was on the other end of the line (calling from Paris, I believe). In a professional manner I was guided through a work-around to my mouse problem. Excellent service!

With DrakX up and running I was told to switch to the console (with Ctrl+Alt+F2). From there I could issue the following commands (in dir /dev):

```
mknod psaux c 10 1
ln -sf psaux mouse
```

First "mknod" is used to create a special file for the PS/2 mouse device, and then a symbolic link is made to point from psaux to mouse.

It must be my specific system which is a little picky, because I personally haven't heard of other users having the same kind of mouse problem. I'm sure MandrakeSoft with future releases of their distribution will come up with a little fix for this potential mouse problem.

Well, let's get on with the installation... I jump back in DrakX with Alt+F7, and my mouse is now back from the dead, but it's not responding correctly to movements, so I'll stick to my keyboard for the next few steps, until I get to choose my mouse type a little later during the install. First you get to choose language for

the install procedure, and then you have to accept the license agreement.

Next comes choice of installation class (recommended, customized, or expert). Choosing 'recommended' is probably the choice of a first-time Linux installer, but 'expert' gives you a lot more control over the install procedure, and 'expert' isn't really all that difficult.

I don't see a need for the middle-way choice called 'customized'. 'Expert class' shouldn't be a problem at all for people who have installed Linux just a couple of times, although not yet being any expert at it. DrakX is simply just too user-friendly for the 'expert class' to be very difficult to use. This is meant as a compliment.

After choosing 'expert' you need to specify what the system mainly will be used for (either as a workstation, server, or as a development machine). Let's go with the workstation install for this review as most people probably want to do this.

Finally, I get to choose my PS/2 mouse from a mouse list, which makes it work just the way it should!

A few steps later I choose security level for my system ('medium' is probably okay for most people), and I choose to use Supermount (easy mounting of floppy, cd-rom, etc - like in Windows). It amazed me that Mandrake still seems to be the only distribution on the planet to include Supermount as an option during install. I honestly cannot live without it. Thumbs up to Mandrake for this one.

I now get to partitioning and formatting my hard drive. The graphical tool included with DrakX is very easy to use and navigate (if you have a working mouse, that is ;)). Formatting partitions took a very long time. I'd like to see it go a little faster if possible. Other distributions (e.g. Best Linux) do it a lot faster (at least on my system), so it can be done.

Package selection is easy with DrakX and dependencies can be solved automatically when choosing between individual packages. A full install takes up about 3.5 GB of space on your hard drive, but you can, of course, choose a much smaller amount to install, if you wish.

Besides given the opportunity to configure your network settings, there are configuration options for setting up a connection to the Internet using either ISDN, xDSL, PPP (modem), or a cable connection.

When coming to the step that allows you to choose which services should be automatically started at boot time, I suggest that you deselect all the services you won't be needing. If your intentions are to use your system as a workstation, you probably don't want a web server, database server, mail server, and a lot of other server services.

I have never really understood why most distributions seem to think that a regular desktop/workstation user wants all these memory consuming server services to be running in the background for no reason. Some of them are also rather insecure when not properly post-configured by the user.

When it comes to printing, DrakX offers you to choose between the well-known lpr or the new, and more powerful, CUPS (Common Unix Printing System). You can setup local, remote and SMB printers.

I am very pleased to see that - for the first time ever - a distribution is able to autodetect my cheap Canon BJC-2000, but for some reason printing a test page didn't work. Had the same problem with Mandrake 7.1 (and also Best Linux 2000 R2), but the printer always works regardless of this failed test page print.

Configuring X works okay, but it could (and should) be better. My old TNT is for some reason detected as a Riva128, so I have to manually choose my TNT. Not a big thing, but it should not be an issue with today's distributions. With my TNT I get to choose between XFree86 3.3.6 with 3D hardware acceleration (experimental), XFree86 3.3.6 (without 3D accel.) and XFree 4.0.1.

Even though DrakX offers me 3D hardware acceleration out-of-the-box with 3.3.6 (which is a nice feature, indeed) I go with the newer 4.0.1 in hope for better all-round performance. I can always include support for 3D hardware acceleration later on myself.

My monitor (a Samsung SyncMaster 1100p+) is also not detected, and it cannot be chosen manually from the list of monitors, so I have to go with the "Multi--frequency monitor that can do 1280x1024 @ 74Hz" which is a pretty close match to my monitor specs. (it can do 1280x1024 @ 75 Hz).

Although I've never tried it out myself, it seems like a really nice feature that DrakX offers you to generate an auto install floppy for Linux replication at the end of the installation. This will come in handy, if you have multiple machines with the same hardware configuration.

After completing the installation procedure the system reboots. I am very surprised to see that Mandrake 7.2 comes with a new graphical boot. I like it somewhat, but I'm sure some Linux users (of the "old school") will say it is a bit too colorful. Well, it is all in the eye of the beholder.

RUNNING X

When I launch X I'm a bit disappointed to see that my 1280x1024 default screen resolution is only running at 60 Hz. Why not 74 Hz, like I specified during install? I ran xf86config to manually re-configure X which is rather easy, if you know your system specs., but for a newbie it's not a funny thing to be using a tool like xf86config.

Using xf86cfg might be a better option for some, but I have a feeling that this graphical tool really doesn't do the job well enough for most people. At least it did not for me when I tried it.

KDE 2

With KDE 2 as the default desktop environment I'm feeling right at home. Although the KDE shipping with Mandrake 7.2 isn't the KDE 2 FINAL, it is darn close to being. If I remember correctly, it is the RC3 (Release Candidate 3) which is only about one week away from being KDE 2 FINAL. Nevertheless, if you feel like upgrading (KDE 2.0.1 is out already), just visit this place and grab the RPMs.

Playing an MP3 from XMMS tells me that my SB Live! is properly configured, which I appreciate. Unfortunately, I don't have my CD recorder at the moment, so I won't be able to test if my old Philips CDD 3610 works right out-of-the-box with Mandrake 7.2. My guess is that it would have worked just fine, just like it did with my good ol' Mandrake 7.1.

Mandrake 7.1 was the very first distribution ever to being able to pre-configure my CD recorder. I've always liked Mandrake for its ability to detect and properly configure the more "exotic" hardware such as CD recorders. My experience with different distributions tells me that no other distribution can match Mandrake at this point, although SuSE comes pretty close.

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

http://www.linux-world.dk/reviews/mdk72-review.php

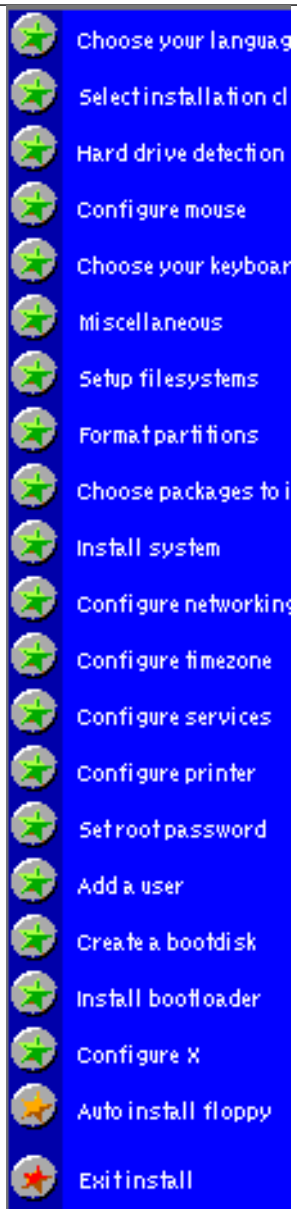

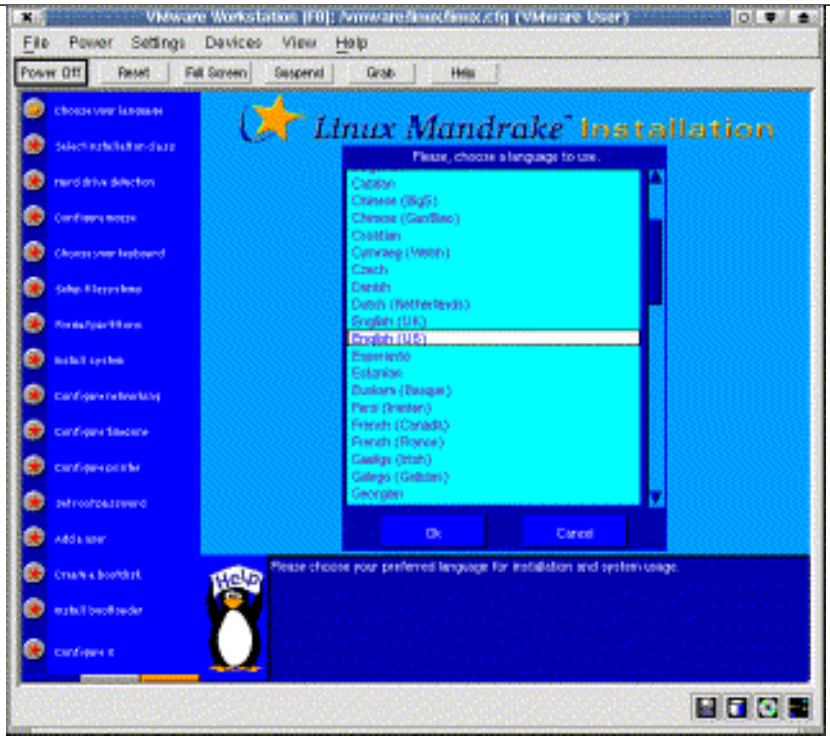

Language

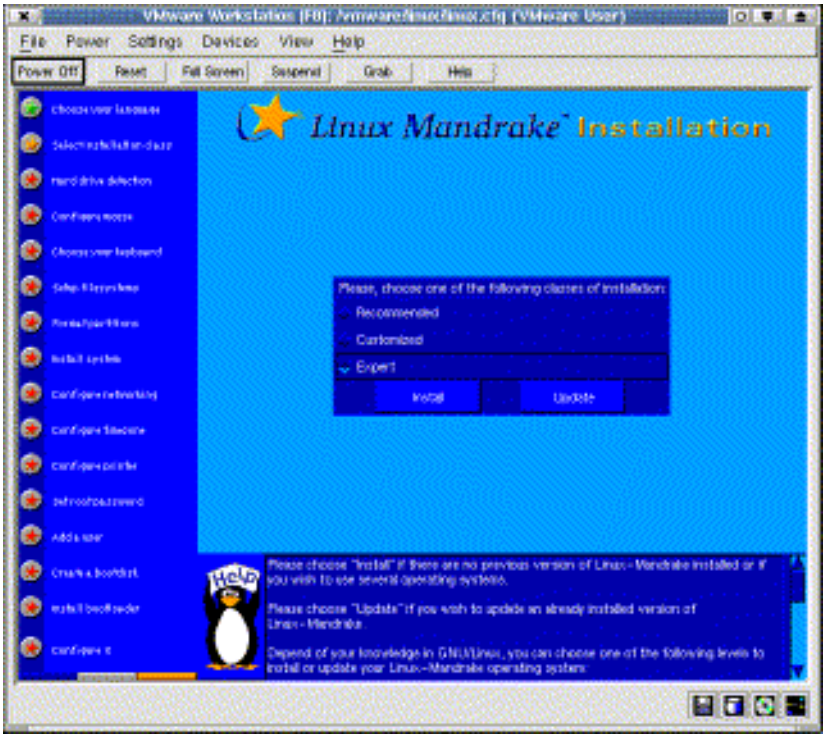

Installtion Type

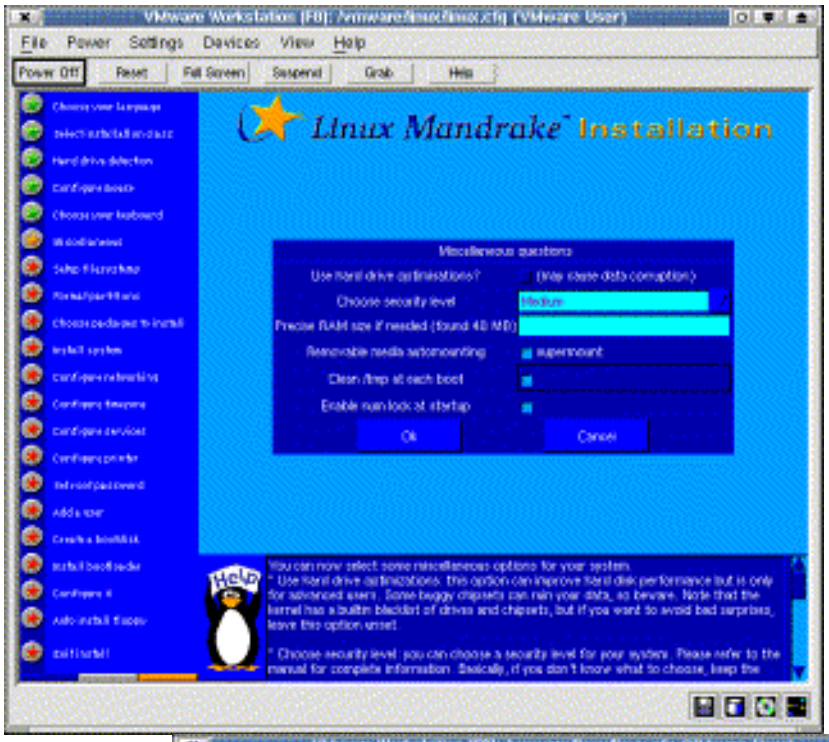

Security Levels

Disk Partitioning

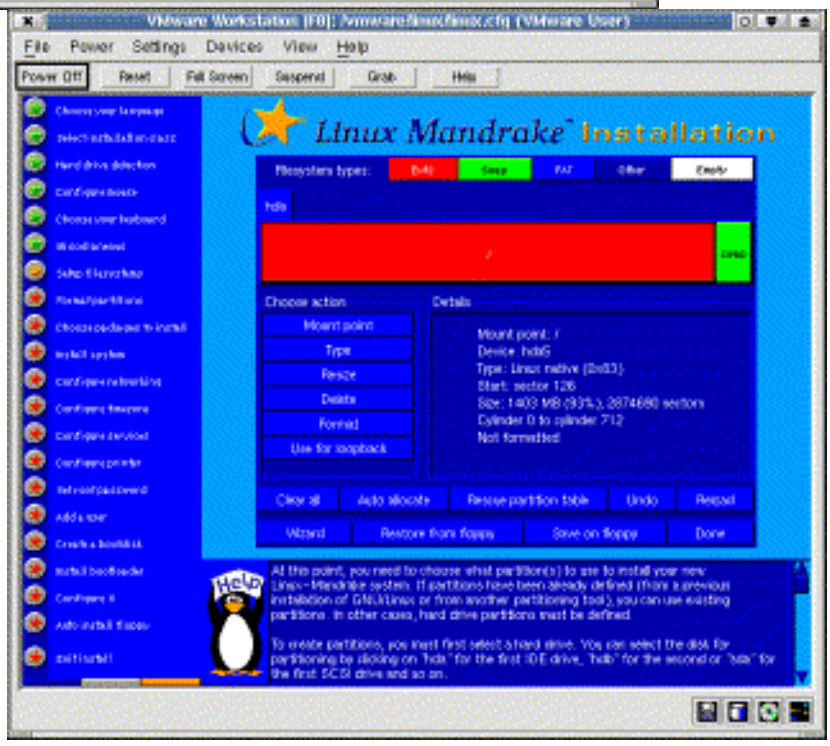

Firewall Security

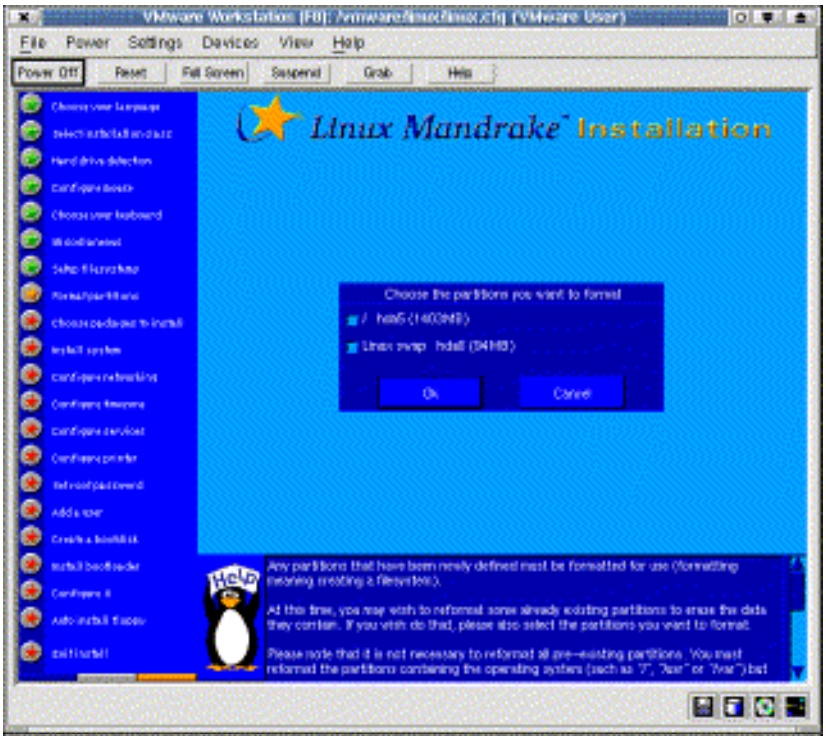

Package Selection

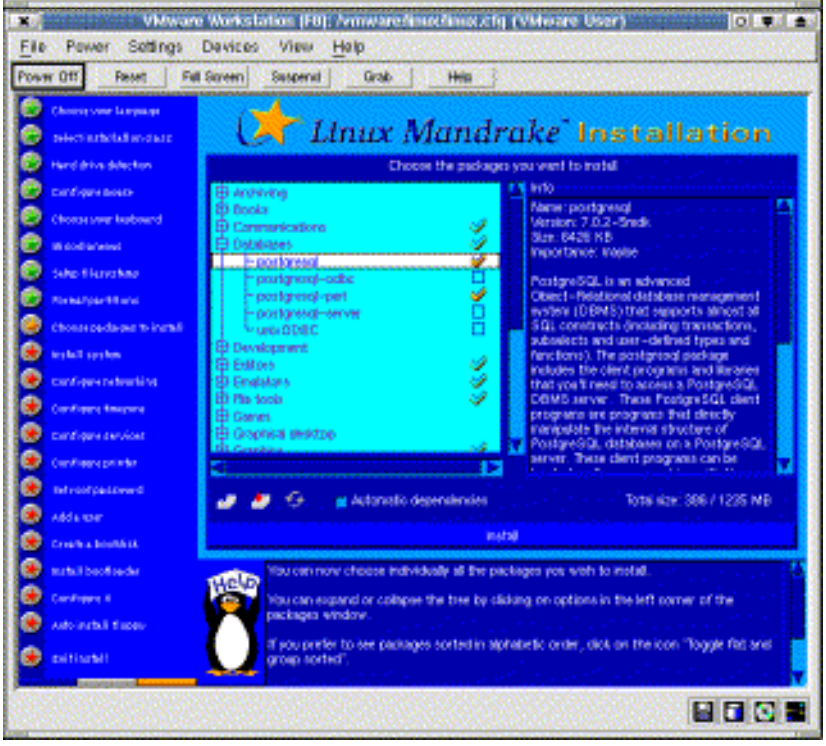

Installing now

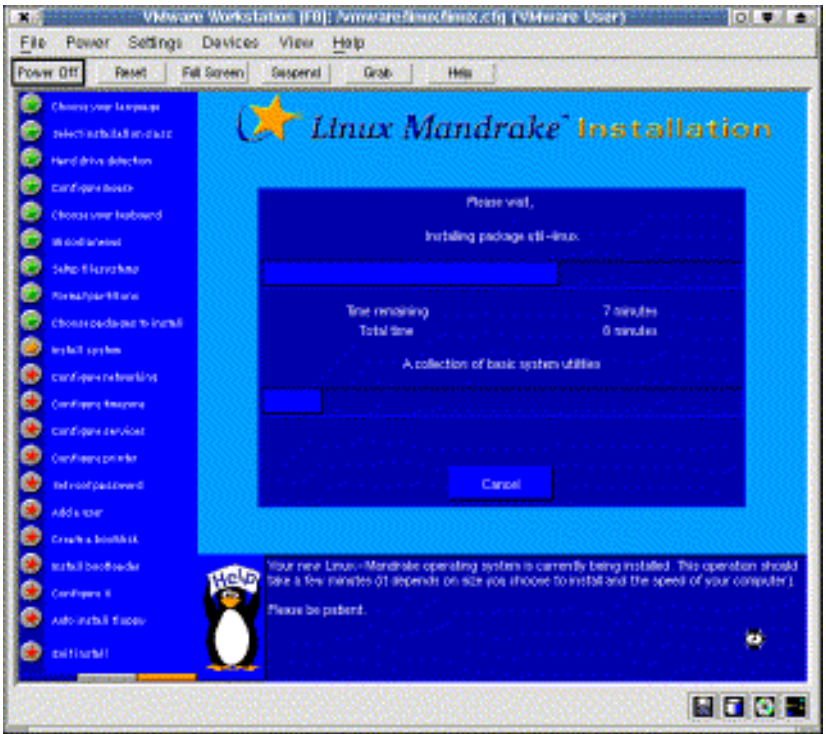

Network Config

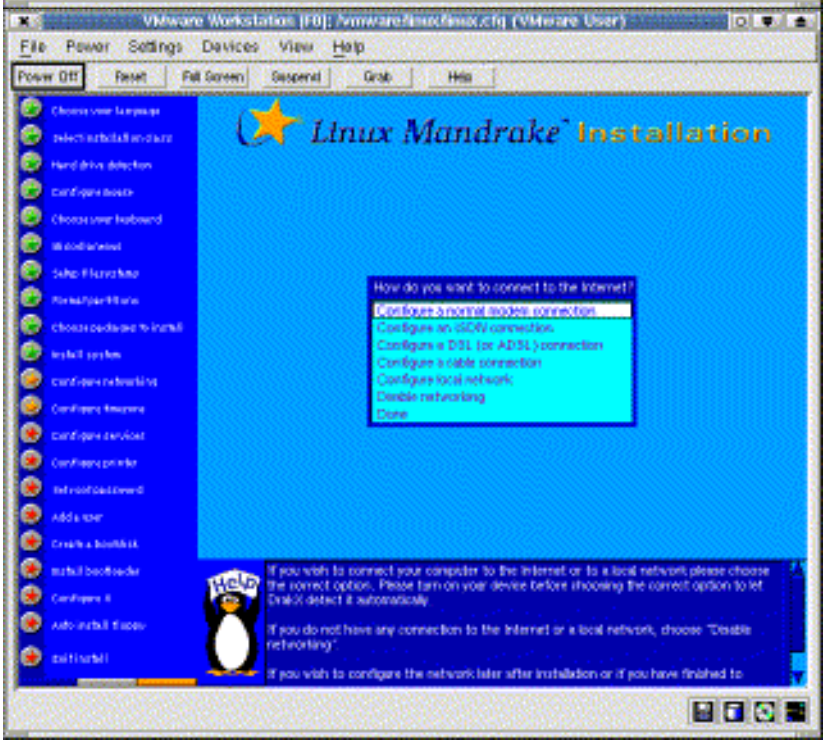

Configuring Unix Service

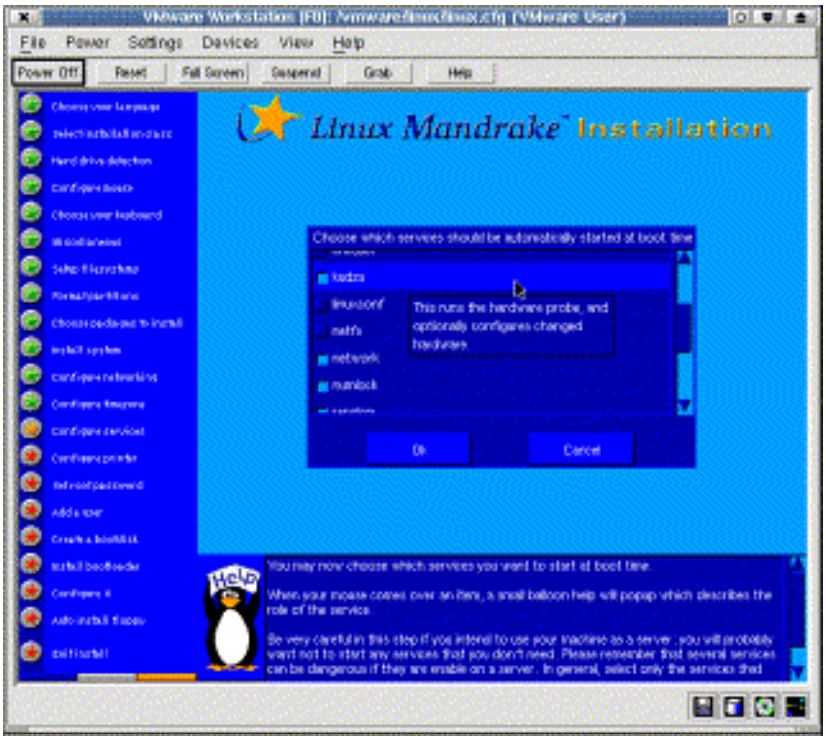

Printing Config

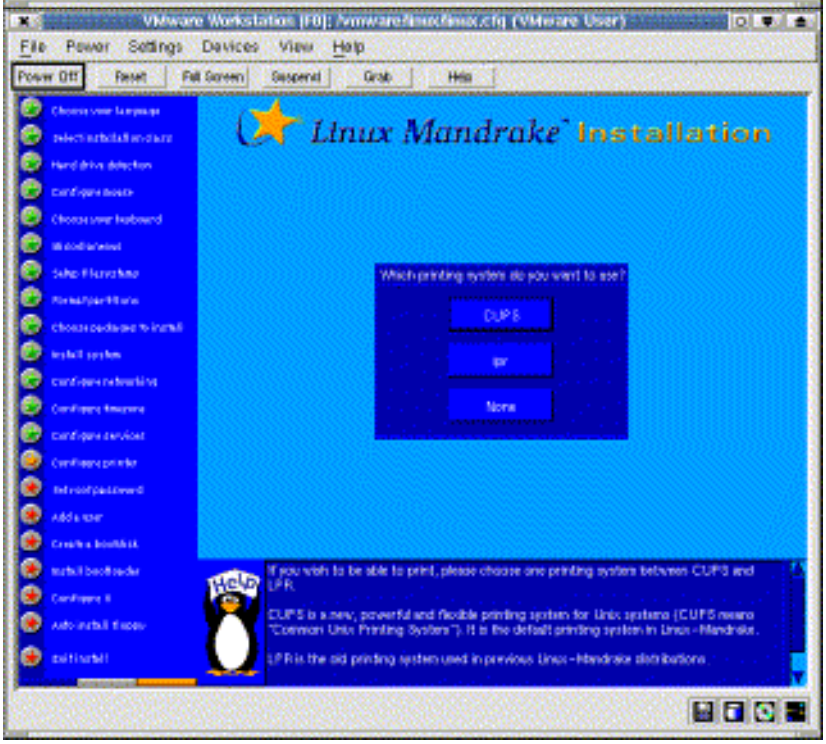

Choosing the Printer

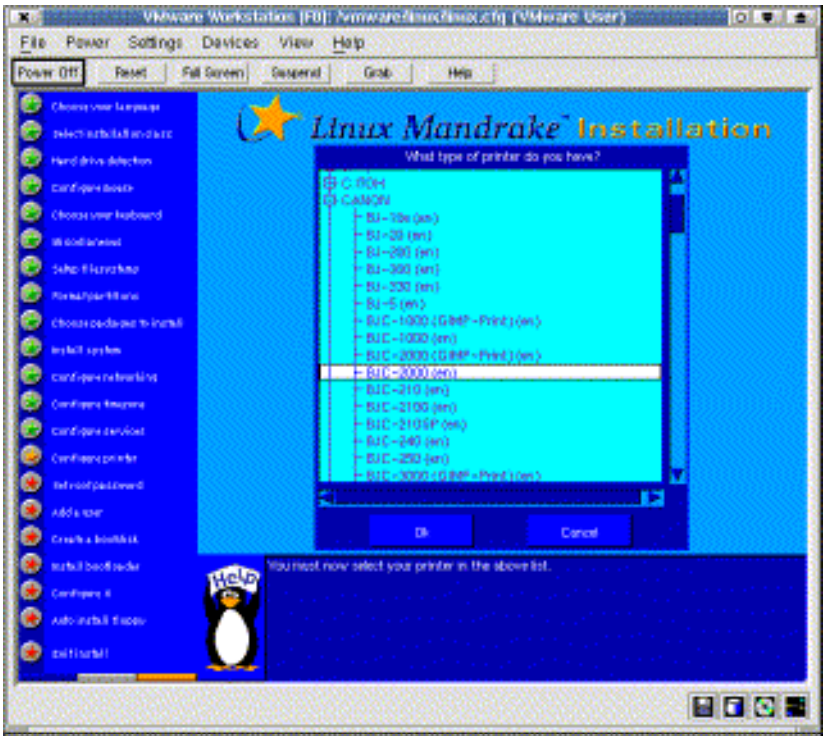

Graphics Card Configuration

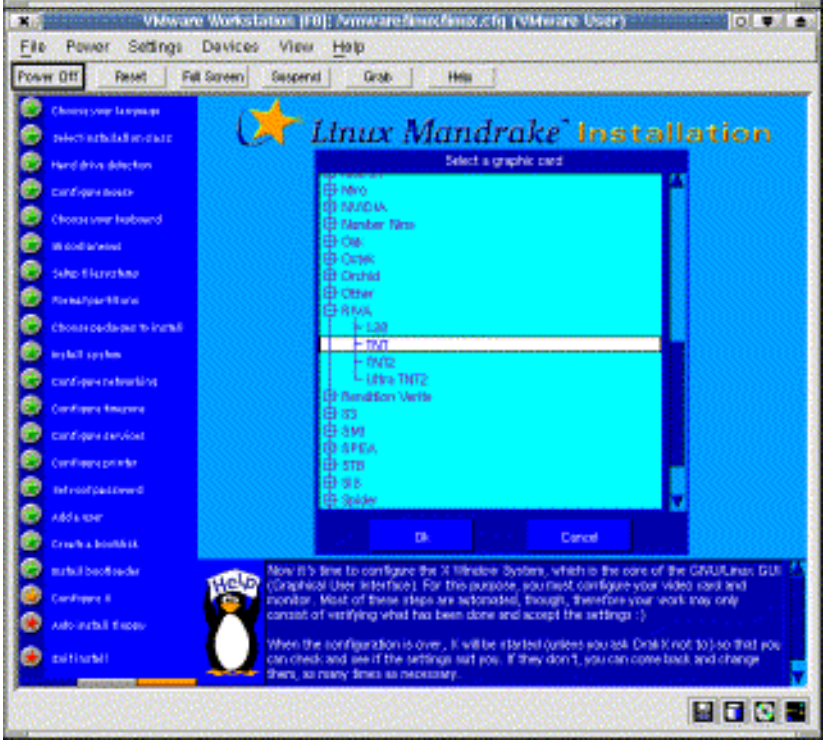

XFree86 Selection

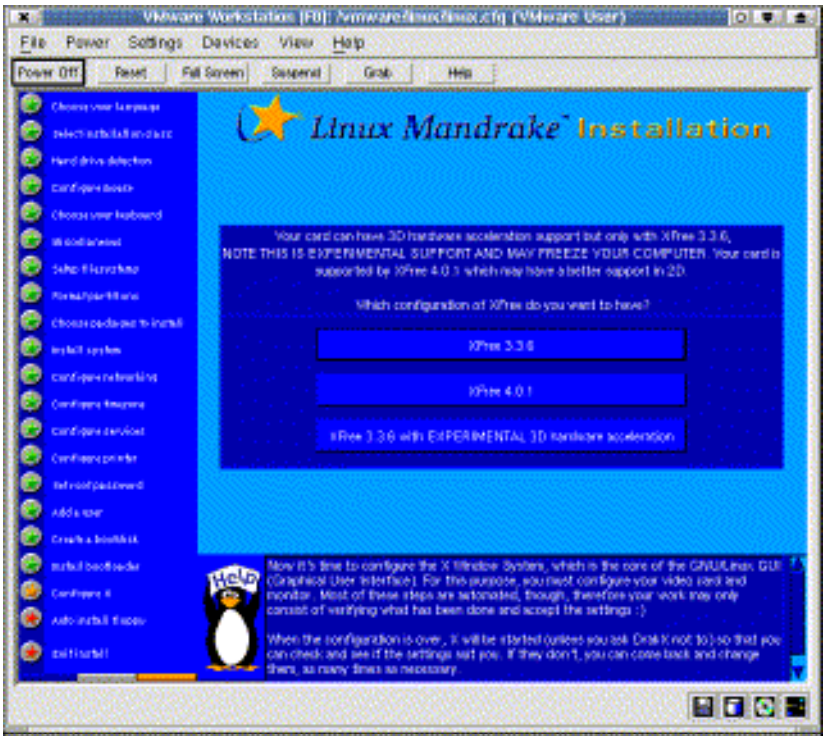

Monitor Selection

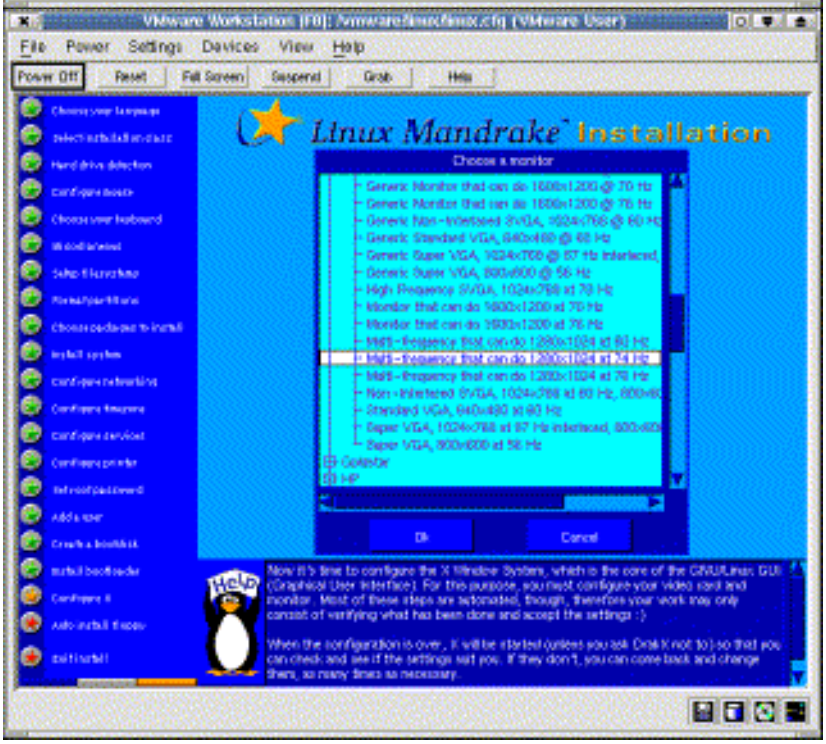

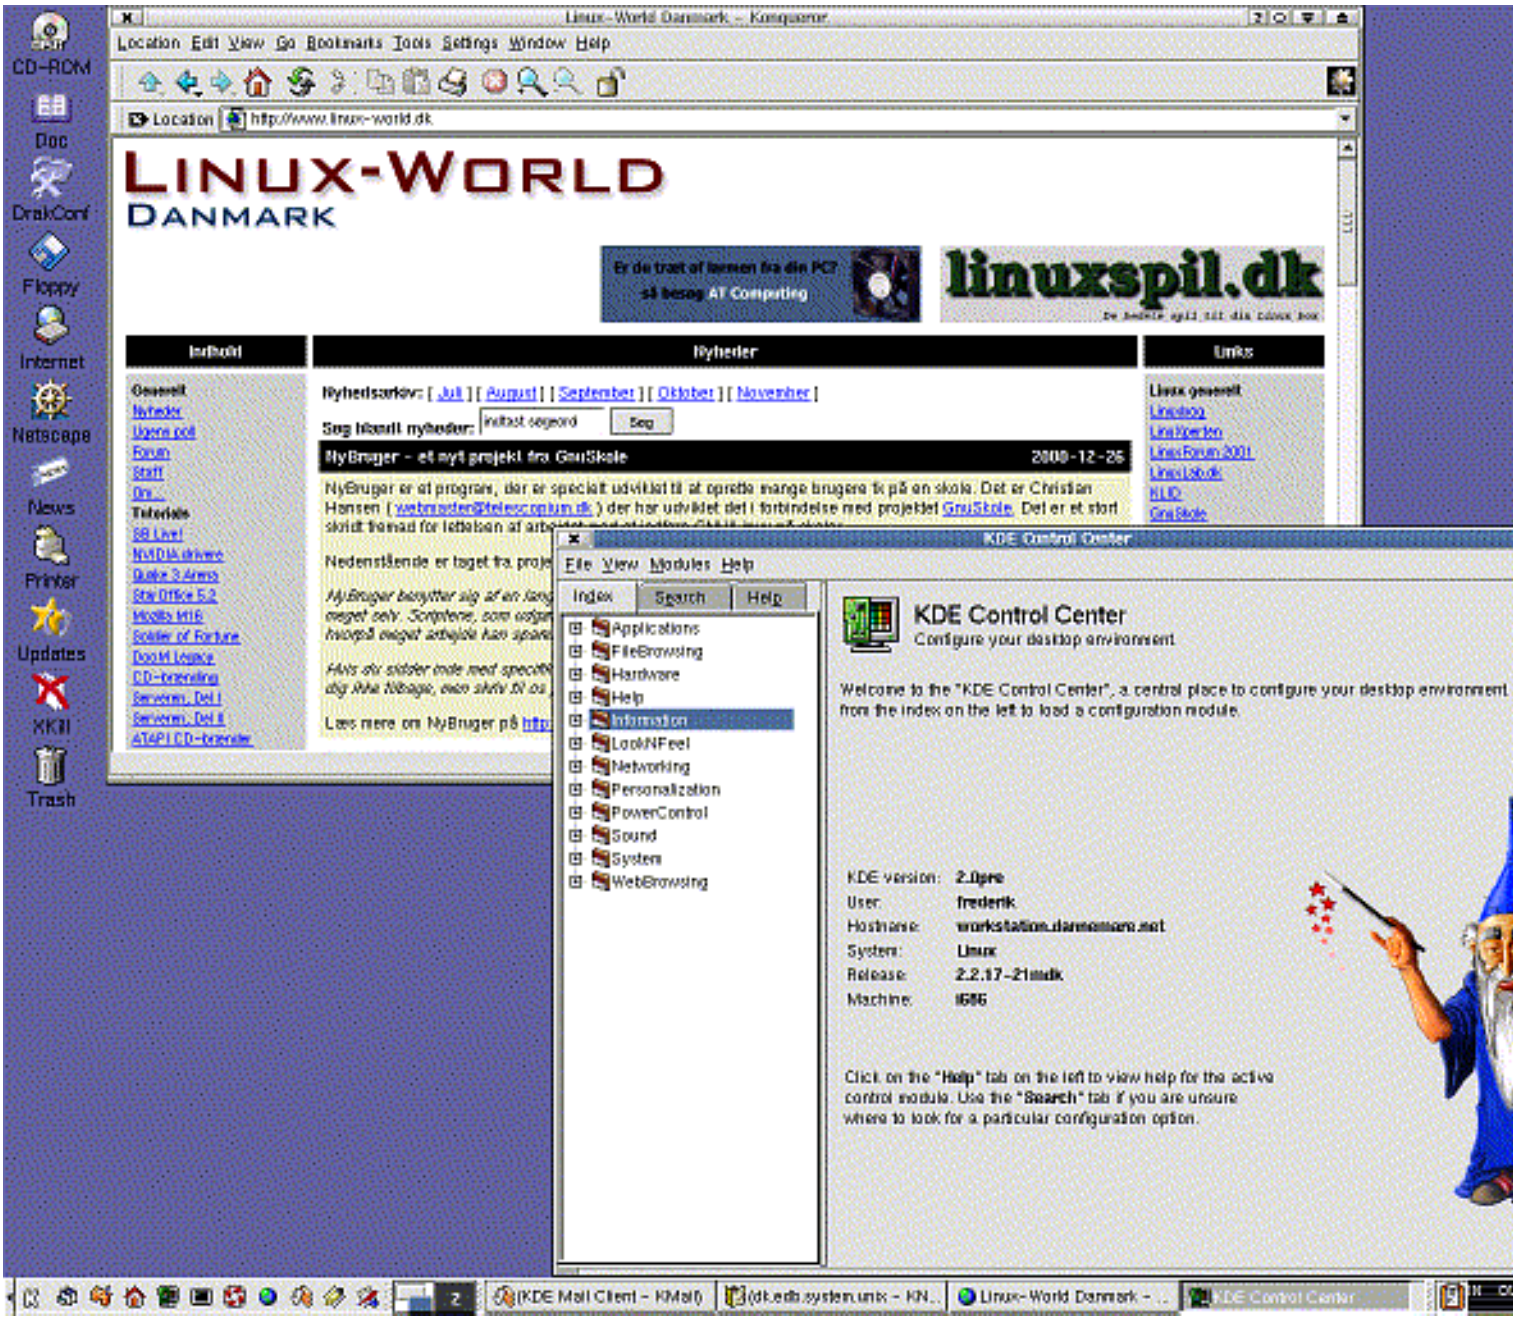
CD-ROM
Doc
DrakConf
Floppy
Internet
Netscape
News
Printer
Updates
XKill
Trash
Linux-World Danmark - Konqueror
Location Edit View Go Bookmarks Tools Settings Window Help
Location http://www.linux-world.dk
LINUX-WORLD
DANMARK
linuxspil.dk
Indhold
Nyheder
Links
Nyhedsarkiv: [Juli] [August] [September] [Oktober] [November]
Søg tilbage til nyheder:
Søg
NyBruger - et nyt projekt fra GnuSkole
2000-12-26
KDE Control Center
File View Modules Help
Index Search Help
Applications
FileBrowsing
Hardware
Help
Information
LookNFeel
Networking
Personalization
PowerControl
Sound
System
WebBrowsing
KDE Control Center
Configure your desktop environment
Welcome to the "KDE Control Center", a central place to configure your desktop environment from the index on the left to load a configuration module.
KDE version: 2.0pre
User: frederik
Hostname: workstation.danmosers.net
System: Linux
Release: 2.2.17-21mdk
Machine: i686
Click on the "Help" tab on the left to view help for the active control module. Use the "Search" tab if you are unsure where to look for a particular configuration option.
(KDE Mail Client - KMail)
Linux-World Danmark - ...
KDE Control Center

The Open Source Lucky Dip

By: Con Zymaris conz@cyber.com.au

Welcome back.

A bit of a bumper crop of apps and utils for you this ti me around. As always, you know where to send pointers to juicy apps that you want listed here, right? auugn@auug.org.au

###

Humorix: Australia To Ban Windows?

Written by James Baughn on July 8, 2001
from the a-law-that-actually-serves-the-public-good dept.

Two years ago, Australia passed legislation that effectively outlawed Linux. Now, the nation appears poised to pass legislation that would effectively throw all Windows users in jail.

Australia wants to make it illegal to possess "hacking paraphernalia", including the source code to malicious viruses. However, every Windows computer connected to the Internet has undoubtedly received at least one macro virus coded in Visual Basic.

Thus, millions of users have the source code to a virus sitting somewhere on their hard drives.

Thus, thousands of email servers have the source code to a virus sitting in some poor schmuck's inbox.

Thus, hundreds of Internet Service Providers have the source code to a virus flowing through their network even as you read this sentence.

The bottom line is that countless Australians could shortly become criminals because they use Windows and Outlook. Oh darn.

The developer responsible for integrating Visual Basic into Microsoft Outlook was unavailable for comment at press time.

http://i-want-a-website.com/about-linux/jul01.shtml#Windows-Ban

###

Dialog

Dialogm by developer Vincent Stemen, lets you to present a variety of questions or display messages using dialog boxes from a shell script (or any scripting language). These types of dialog boxes are implemented: yes/no box,menu box, input box, message box, text box, info box, guage box, checklist box, file-selection box, and radiolist box. Dialog is GPL. Get it from: http://www.AdvancedResearch.org/dialog/

AutoDia

If you're a developer, check this out. AutoDia is a Perl application designed to allow the easy creation of XML diagrams from various data sources. The output is meant for use with Dia (or any XML parser). AutoDia supports any language through the use of handlers, and a good handler for Perl as well as a simple handler for C++ are available. Download it from
http://droogs.org/autodia/

AWOL

Are you trying to make your work envirnoment more organised? Is trying to organise geeks like trying to herd cats? ;-) Ok, AWOL is an in/out board, similar to something you'd see on the wall of a lobby or desk of a switchboard operator, to keep track of who's in the building, when they'll be back, etc. It's accessible from the Web, so people can mark themselves in or out from their own desk. AWOL isn't AWOL at:
http://gospelcom.net/mnn/topher/awol/

All Work and No Play: Dave Gnukem

Written by David Jaffe, Dave Gnukem is a 2D scrolling platform game, similar to Duke Nukem 1. It includes a sprite and level editor. It uses GGI and thus runs on the console as well as in X, windowed or fullscreen. It also runs on Windows, using DirectX. Download it and blow yourself away:
http://www.scorpioncity.com/djgame.html

db4o: A 50kB Java object database.

Now this is something that many coding shops I know about would kill to have. Carl Rosenberger has written a cool tool called db4o (database for objects), which is a Java object database. It features automatic object recognition, query-by-example interface, and navigation. Changes to application classes are not necessary. It's free for non-commercial use from here:
http://www.db4o.com/

Distributed Checksum Clearinghouse

Tired of Spam? Tired of inane questions? ;-) Then perhaps add this into your systems mix: Distributed Checksum Clearinghouse (DCC) is a system of clients and servers that collect and count checksums related to mail messages. The counts can be used by SMTP servers and mail user agents to detect and reject bulk mail. DCC servers can exchange common checksums, which checksums include values that are constant across common variations in bulk messages. Written by Vernon Schryver, DCC is available from:
http://www.rhyolite.com/dcc/

ext3

What in the end may become the de-facto journaling d file-system for Linux, ext3 extends the current ext2fs is presently in heavy development (primarily by Stephen Tweedie) and is available from:
http://beta.redhat.com/index.cgi?action=ext3

FreeLords: Embrace, Extend Extinguish!

FreeLords is a turn-based strategy game similiar to W*rl*rds. It can be played with friends on one computer or via a network. Written by Michael Bartl, it's available from:
http://www.freelords.org/

ICQ-IRC Gateway

Are you drowning in Instant Messaging systems? Do you prefer the open IRC protocols? Check this out. The IRC-to-ICQ Gateway bot provides a fully functional method to allow messages sent to the bot to reach users online on the ICQ network. This has several advantages. First, it allows IRC users to communicate with ICQ users without requiring any ICQ software. It can also be used to leave short messages for ICQ users with more speed and ease than email. Written by Stanislav Lechev, it is available from:
http://firedust.vega.bg/icqserv.html

Java Remote Method Invocation Language

There has been a lot of noise recently about remotely executable content on the Web. While noting new to people accustomed to good ol' RPC, some of the ideas and preocesses are interesting. jRML (Java Remote Method Invocation Language) is a Java object (de)serialization and messaging library which uses a more compact format than XML-RPC. According to the author, Bernhard Fastenrath, it is also easier to use than most SOAP implementations. More information from: http://xrml.sourceforge.net/

KMyMoney2

Need a Quicken-like app on KDE? Try this. KMyMoney2 is a KDE personal finance program. It's currently in pre-release form, and available from:
http://kmymoney2.sourceforge.net/

Yngwie: Be your own Movie Mogul ;-)

Yngwie is a graphical frontend for mpeg2divx and libcss to help you RIP movies off DVD disks and compress them into DivX format. Hmmm, *fair use* ;-)
Get it from: http://yngwie.naken.cc/

wxDesigner

Last issue, we looked at wxPython. Time to look at a related project. wxDesigner is a dialog editor and RAD tool for the wxWindows C++ library and its popular Python and Perl bindings. It includes a visual dialog editor, a bitmap editor, a syntax-highlighting source editor, and built-in mechanisms for automatic generation of file skeletons, GUI classes, event handlers, and getter functions. It provides an identical user interface and identical functionality for C++, Perl, and Python, and it can generate output in C++, Python, Perl, and XML. Download from: http://www.roebling.de/

Scientific Image Database

SIDB (Scientific Image Database) is a Web-driven database for (scientific) images. Entry of image meta--data is facilitated through the use of user-definable templates. Users have complete control over who else may see the submitted image and data. SIDB offers various views on the actual image data. Thumbnail plus meta-data can be printed, and multiple thumbnails can be combined in galleries. A freeware version of the Huygens software (http://www.svi.nl) can be used to generate projections from 3-D images (most confocal microscope image formats are supported), as well as MPEG movies, showing the individual layers from the 3-D images.

###

Microsoft Bundles Worm with IIS

Redmond, WA - Microsoft announced that it bundled a worm with its latest version of Internet Information Server.

"We did it because it's beneficial to our users," said Bill Gates. "With the worm already bundled as part of the software, network administrators won't have to wonder if their system is infected. They will know. It's obvious with the number of servers already infected with worms and Trojan horses that this is something our users want. Otherwise, why would they be infected?" ...

http://www.bbspot.com/News/2001/07/worm.html

###

Advertisement: American Bookstore

Getting Out of MS Access [to MySQL]

Author: Cory R. Rauch <osfaqcom@osfaq.com>

You just finishing converting your servers to Linux, but now want to get the corporate MS Access database into a MySQL database. How do you do that? Well, In this short tutorial we are going to show how to convert an MS Access database into a MySQL database. To accomplish this task we are going to use a script written by Brian Andrews, that converts Access tables into MySQL tables. So lets get started.

The Conversion

First you need to download the script. You can find it at the below link:

http://www.osfaq.com/downloads/a2sql.txt

Next, We open the database we want to convert in Access and cut and paste this script into the module section.

Then we define a new macro with the RunCode () function and call the export_mysql () function.

Finally we run the macro we defined above. The script should rip through your databases and create a mysqldump.txt file. This file contains the SQL statements necessary to re-construct your database MySQL.

Loading Data into MySQL

On the Linux side we need to load the mysqldump.txt file into MySQL. To do this simply type the following at the prompt:

```
# mysql -u root mydb < mysqldump.txt
```

Note: Substitute the mydb with your database alias and you probably need to specify a password too when running the mysqldump.txt file

If everything went ok, then you should have your companies database on Linux/MySQL.

Conclusion

I hope you found this short tutorial useful. Now if only someone could write a setup program that converts all your data and installs Linux in one shot. Please check back for our next article on Linux.

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

http://www.osfaq.com/article.php3?sid=76

The CD-ROMs in AUUGN

This is the third issue of AUUGN which includes complimentary CD-ROMs. We started in January 2001 with the complete Red Hat 7.0 distribution, and followed up in April with FreeBSD 4.2. After that issue, I sent out a mail message to AUUG members asking for their opinion of the service.

The results were overwhelming. A good 10% of the membership replied. That may not seem much, but I think it's the biggest reply quota we've ever had *** check this ***. Even more interesting was that, without exception, all the answers were positive.

One question I asked was what you would like to see in future issues of AUUGN. Here's a breakdown of the useful answers (multiple choices were possible):

OpenBSD	13
NetBSD	5
FreeBSD	3
BSD (unspecified)	3
Total BSD:	24
Linux (unspecified)	5
Debian GNU/Linux	5
Red Hat Linux	3
Suse Linux	3
Mandrake Linux	2
Slackware Linux	1
Debian GNU/Hurd	1
Total Linux:	20
StarOffice	9
Darwin (MacOS X kernel)	3
Solaris X86	2
AUUGN on CD	1
Plan 9	1

Other suggestions were for Linux PPC and SPARC, and Microsoft NT (for drink coasters).

The Linux distributions are a problem, expecially Debian, the most popular: the number of CDs is simply too much to send easily with the AUUGN. By far the most popular one, however, is OpenBSD, so that's what we're distributing with this issue. See the article on OpenBSD elsewhere in this issue, and please note that we're soliciting donations for OpenBSD, which draws most of its revenues from CD sales.

For the next edition, we're investigating the second favourite, StarOffice. Based on the feedback we've been getting, we'll also try NetBSD and another Linux distribution in the issues to come. We're still open to suggestions, so if there's something else you'd like, please let us know.

AUUG Security Symposium

19-21 November 2001 - Brisbane
Call for Papers/Presentations and Tutorials
AUUG Security Symposium 2001
Brisbane

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 practitioners 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 un-ashamedly for the practitioner in the field who wants to share or learn more about how to secure their systems (be it a PC operating systems, a huge network or a client server application).

The AUUG security symposium will be a three day event with paper presentations and tutorials. The symposium is to be held from Monday, 19 to Wednesday 21 November 2001, at the Auditorium of the Primary Industries Department, 80 Ann St, Brisbane. Monday will be for tutorials and the Tuesday and Wednesday for paper presentations.

Call for presentations/papers and tutorials:

Authors are invited to submit abstracts associated with information security presenting any of the following:

- new or interesting results of current research,
- insight into the security of various technologies, products or systems
- recent development work,
- implementation or deployment approaches/techniques,
- security management lessons/experiences. This includes submissions on any topic relating to technical security issues of UNIX, Linux/FreeNIX, open source, development, networking, and open-systems in the widest sense.

Speaker Incentives:

Paper presenters are afforded free conference registration. Potential tutorial presenters should contact the conference committee regards options for remuneration.

Form of Submissions:

Please submit an abstract together with an outline. The outline should contain enough detail to allow the program committee to make a reasoned decision about the final presentation. An abstract should include:

1. Author names(s), postal addresses, telephone numbers, fax and e-mail addresses.
2. A biographical sketch not to exceed 100 words.
3. Abstract: 100 words.
4. Audio-visual requirements: Please indicate your requirements for overhead projector or video/computer equipment.

Acceptance:

Authors whose submissions are accepted are asked to provide a presentation in some machine readable format which can be converted to HTML or PDF. A formal paper is not required, but is welcomed. Copies of presentations will be made available on the WWW after the conference.

Relevant Dates:

Abstract and outlines due: 31 August 2001
Acceptance: 30 September 2001

Exhibitors:

There will be a exhibition area as part of the conference for product vendors and service/consultancy providers. Expressions of interest can be addressed to Liz Carroll (AUUG Business Manager) on busmgr@auug.org.au or phone 02 8824 9511 or 1800 625 655 or Fax (02) 8824 9522

Sponsorships Opportunities:

Sponsorships are available. For further information, contact Liz Carroll (AUUG Business Manager) on busmgr@auug.org.au or phone 02 8824 9511 or 1800 625 655 or Fax (02) 8824 9522

Addresses:

Please submit hard copy or electronic copy (preferred) of abstracts and outlines to:
g.gaskell@bigpond.com

Organising Committee:

Gary Gaskell - Bank of Queensland
Warren Toomey - Bond University
Duncan Unwin - QSI
Lawrie Brown - ADFA
Liz Carroll - AUUG

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.auor phone 02 8824 9511 or 1800 625 655 or Fax (02) 8824 9522.

Meet the exec: Greg Lehey

I was born in Melbourne, but I left Australia at the age of 5 and spent most of my life overseas. I went to school in Malaysia and England and studied Chemistry at the University of Hamburg in Germany and Chemical Engineering at Exeter in England, thus ensuring that I have no qualifications in the area of Computer Science.

After University I returned to Germany in search of a girl I never found, and spent twenty-five years there working in just about all areas of the profession: with computer manufacturers such as Univac, Tandem, and Siemens-Nixdorf, the German space research agency DFVLR, nameless software houses and a large end-user. For the last eight years I worked for myself. In the process I became very involved in the open source community, specifically BSD UNIX, which I've been using as my main operating system since 1992.

I consider myself a jack-of-all-trades rather than a specialist:

I've performed most jobs, ranging from kernel development to product management, from systems programming to systems administration, from processing satellite data to programming petrol pumps, from mastering and producing CD-ROMs to DSP instruction set design. About the only thing I haven't done is writing commercial applications software.

I returned to Australia in early 1997 and live on a 50 acre farm in the Adelaide Hills, not quite beyond the Black Stump (in fact, it's in a paddock in the South--East of the property). I work from home for IBM Australia in the Ozlabs division of the Linux Technology Centre. A self-professed non-user of Linux, I spend much of my time trying against all odds to persuade IBM to switch to BSD.

I am also an active writer. I have written 3 books: "Porting UNIX Software" (O'Reilly and Associates, 1995), "Installing and Running FreeBSD" (Walnut Creek, Concord CA, 1996), and "The Complete FreeBSD" (BSDi, 1999). In addition, I write a number of magazine articles, notably in Daemon News (http://www.daemonnews.org/), where I have been writing a bi-monthly column called "The Daemon's Advocate" since October 1998, alternating with Wes Peters, who does the other months.

One of the central themes of these columns is the cooperation between various free software projects. I also keep a daily diary on the web at http://www.lemis.com/grog/diary.html.

I've been on the AUUG executive committee since August last year, and since 1 July 2001 I have been the secretary. I am also a member of the FreeBSD Core team and the secretary of the Peruvian Horse Registry of Australasia. You can reach me by mail at Greg.Lehey@auug.auug.au, grog@FreeBSD.ORG or grog@lemis.com.

There's no difference between these addresses: all mail ends up at the last address. Alternatively, browse my home page at http://www.lemis.com/grog/.

I consider myself a perfectionist (that's my interpretation of the more typical claim "bloody nitpicker"). I have gradually learnt to tone down my objections to what I view as mediocrity, and I now no longer annoy people all the time. I'm a strong believer in doing things as correctly as possible, and have occasional run-ins with the other committee members when it comes to issues like the use of Microsoft products.

I still try to have a life. When I can drag myself away from my shed full of UNIX workstations, I am involved in performing baroque and classical woodwind music on my collection of original instruments, exploring the South Australian countryside with my family on our Arabian and Peruvian horses, or exploring new cookery techniques or ancient and obscure European languages.

IPtables Tutorial

By: David LeCount <snailboy1@yahoo.com>

I'm sure many of you have been wondering how to use iptables to set up a basic firewall. I was wondering the same thing for a long time until I recently figured it out. I'll try to explain the basics to at least get you started.

First you need to know how the firewall treats packets leaving, entering, or passing through your computer. Basically there is a chain for each of these. Any packet entering your computer goes through the INPUT chain. Any packet that your computer sends out to the network goes through the OUTPUT chain. Any packet that your computer picks up on one network and sends to another goes through the FORWARD chain. The chains are half of the logic behind iptables themselves.

Now the way that iptables works is that you set up certain rules in each of these chains that decide what happens to packets of data that pass through them. For instance, if your computer was to send out a packet to www.yahoo.com to request an HTML page, it would first pass through the OUTPUT chain. The kernel would look through the rules in the chain and see if any of them match. The first one that matches will decide the outcome of that packet. If none of the rules match, then the policy of the whole chain will be the final decision maker. Then whatever reply Yahoo! sent back would pass through the INPUT chain. It's no more complicated than that.

Now that we have the basics out of the way, we can start working on putting all this to practical use. There are a lot of different letters to memorize when using iptables and you'll probably have to peek at the man page often to remind yourself of a certain one. Now let's start with manipulation of certain IP

addresses. Suppose you wanted to block all packets coming from 200.200.200.1. First of all, -s is used to specify a source IP or DNS name. So from that, to refer to traffic coming from this address, we'd use this:

```
iptables -s 200.200.200.1
```

But that doesn't tell what to do with the packets. The -j option is used to specify what happens to the packet. The most common three are ACCEPT, DENY, and DROP. Now you can probably figure out what ACCEPT does and it's not what we want. DENY sends a message back that this computer isn't accepting connections. DROP just totally ignores the packet. If we're really suspicious about this certain IP address, we'd probably prefer DROP over DENY. So here is the command with the result:

```
iptables -s 200.200.200.1 -j DROP
```

But the computer still won't understand this. There's one more thing we need to add and that's which chain it goes on. You use -A for this. It just appends the rule to the end of whichever chain you specify. Since we want to keep the computer from talking to us, we'd put it on INPUT. So here's the entire command:

```
iptables -A INPUT -s 200.200.200.1 -j DROP
```

This single command would ignore everything coming from 200.200.200.1 (with exceptions, but we'll get into that later). The order of the options doesn't matter; the -j DROP could go before -s 200.200.200.1. I just like to put the outcome part at the end of the command. Ok, we're now capable of ignoring a certain computer on a network. If you wanted to keep your computer from talking to it, you'd simply change INPUT to OUTPUT and change the -s to -d for destination. Now that's not too hard, is it?

So what if we only wanted to ignore telnet requests from this computer? Well that's not very hard either. You might know that port 23 is for telnet, but you can just use the word telnet if you like. There are at least 3 protocols that can be specified: TCP, UDP, and ICMP. Telnet, like most services, runs on TCP so we're going with it. The -p option specifies the protocol. But TCP doesn't tell it everything; telnet is only a specific protocol used on the larger protocol of TCP. After we specify that the protocol is TCP, we can use --destination-port to denote the port that they're trying to contact us on. Make sure you don't get source and destination ports mixed up. Remember, the client can run on any port, it's the server that will be running the service on port 23. Any time you want to block out a certain service, you'll use --destination-port. The opposite is --source-port in case you need it. So let's put this all together. This should be the command that accomplishes what we want:

```
iptables -A INPUT -s 200.200.200.1 -p tcp --
destination-port telnet -j DROP
```

And there you go. If you wanted to specify a range of IP's, you could use 200.200.200.0/24. This would specify any IP that matched 200.200.200.*. Now it's time to fry some bigger fish. Let's say that, like me, you have a local area network and then you have a connection to the internet. We're going to also say that the LAN is eth0 while the internet connection is called ppp0. Now suppose we wanted to allow telnet to run as a service to computers on the LAN but not on the insecure internet. Well there is an easy way to do this. We can use -i for the input interface and -o for the output interface.

You could always block it on the OUTPUT chain, but we'd rather block it on the INPUT so that the telnet daemon never even sees the request. Therefore we'll use -i. This should set up just the rule:

```
iptables -A INPUT -p tcp --destination-port
telnet -i ppp0 -j DROP
```

So this should close off the port to anyone on the internet yet keept it open to the LAN. Now before we go on to more intense stuff, I'd like to briefly explain other ways to manipulate rules. The -A option appends a rule to the end of the list, meaning any matching rule before it will have say before this one does. If we wanted to put a rule before the end of the chain, we use -I for insert. This will put the rule in a numerical location in the chain. For example, if we wanted to put it at the top of the INPUT chain, we'd use "-I INPUT 1" along with the rest of the command. Just change the 1 to whatever place you want it to be in. Now let's say we wanted to replace whatever rule was already in that location. Just use -R to replace a rule. It has the same syntax as -I and works the same way except that it deletes the rule at that position rather than bumping everything down. And finally, if you just want to delete a rule, use -D. This also has a similar syntax but you can either use a number for the rule or type out all the options that you would if you created the rule. The number method is usually the optimal choice. There are two more simple options to learn though. -L lists all the rules set so far. This is obviously helpful when you forget where you're at. AND -F flushes a certain chain. (It removes all of the rules on the chain.) If you don't specify a chain, it will basically flush everything.

Well let's get a bit more advanced. We know that these packets use a certain protocol, and if that protocol is TCP, then it also uses a certain port. Now you might be compelled to just close all ports to incoming traffic, but remember, after your computer talks to another computer, that computer must talk back. If you close all of your incoming ports, you'll essentially render your connection useless. And for most non-service programs, you can't predict which port they're going to be communicating on. But there's still a way. Whenever two computers are talking over a TCP connection, that connection must first be initialized. This is the job of a SYN packet. A SYN packet simply tells the other computer that it's ready to talk. Now only the computer requesting the service sends a SYN packet. So if you only block incoming SYN packets, it stops other computers from opening services on your computer but doesn't stop you from communicating with them. It roughly makes your computer ignore anything that it didn't speak to first. It's mean but it gets the job done. Well the option for this is --syn after you've specified the TCP protocol. So to make a rule that would block all incoming connections on only the internet:

```
iptables -A INPUT -i ppp0 -p tcp --syn -j
DROP
```

That's a likely rule that you'll be using unless you have a web service running. If you want to leave one port open, for example 80 (HTTP), there's a simple way to do this too. As with many programming languages, an exclamation mark means not. For instance, if you wanted to block all SYN packets on all ports except 80, I believe it would look something like this:

```
iptables -A INPUT -i ppp0 -p tcp --syn --
destination-port ! 80 -j DROP
```

It's somewhat complicated but it's not so hard to comprehend. There's one last thing I'd like to cover and that's changing the policy for a chain. The chains INPUT and OUTPUT are usually set to ACCEPT by default and FORWARD is set to DENY. Well if you want to use this computer as a router, you would probably want to set the FORWARD policy to ACCEPT. How do we do this you ask? Well it's really very simple. All you have to do is use the -P option. Just follow it by the chain name and the new policy and you have it made. To change the FORWARD chain to an ACCEPT policy, we'd do this:

```
iptables -P FORWARD ACCEPT
```

Nothing to it, huh? This is really just the basics of iptables. It should help you set up a limited firewall but there's still a lot more that I couldn't talk about. You can look at the man page "man iptables" to learn more of the options (or refresh your memory when you forget). You can find more advanced documents if you want to learn some of the more advanced features of iptables. At the time of this writing, iptables documents are somewhat rare because the technology is new but they should be springing up soon. Good luck.

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

http://pinehead.com/articles.php?view=371

SMTP over an SSH Tunnel

Jim Mock <jim@freebsdzine.org>

Introduction

When I first got to Massachusetts in February, I was put into temporary housing which was provided to me by work. I was stuck using a dial-up account from home, which wasn't so bad -- at least I had 'net access. However, since I was using Earthlink, and many mail servers block mail originating from them even if relayed through their mail servers (which I was doing), I wanted a way to use a machine at work as my relay instead.

An easy way to do so would have been to allow relay for earthlink.net through the machine I was going to use as my relay. This would have been bad, and as a result was an unacceptable option. You'd have to be a serious cracksmoker to open up a machine to relay for Earthlink. Spammers would sniff it out fairly quickly, and proceed to hammer it. To get around this, I decided to set up an SSH tunnel from a port on my laptop to port 25 on the mail server. It was a fairly easy task, and well worth it in the long run. The remainder of this article will explain how I set things up.

Set up the Tunnel

The first thing I decided was to establish the tunnel as a non-root user. Since the tunnel was going to exist for solely mail relaying purposes, I created a relay user on both my laptop and the server in question. I also ran ssh-keygen(1) and gave the relay user an empty passphrase. If you're overly paranoid, you can use a passphrase and then use ssh-agent(1). The way I figure is if someone gets into my laptop, I have more things to worry about than them sending mail through my relay.

Ok, now that we have the user created, the next step is to actually make sure we can establish the tunnel. This is done by doing the following:

```
% su relay
Password:
% ssh -2 -N -f -L 9595:mail.domain.com:25
mail.domain.com
%
```

Here's an explanation of the above:

The % su relay part allows you to become the relay user. The ssh command is a little more in-depth. The -2 flag tells ssh that we should use SSH v2. We want to use v2 because the ssh process can be sent to the background without running a command using the -N flag. -N is only available in v2 and is very useful when you want to simply forward ports. The -f tells ssh to go into the background. -L tells ssh that the given port on the local side is supposed to be forwarded to the given port on the remote side. Therefore,

```
-L 9595:mail.domain.com:25 forwards local
port 9595 to remote port 25 on mail.domain.com.
```

The mail.domain.com at the end is the hostname of the machine we're connecting to.

After running the command above, you should be able to telnet to port 9595 on the local machine and see that it's really forwarded to port 25 on the remote machine:

```
% telnet localhost 9595
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.
220 mail.domain.com ESMTP Sendmail
8.11.3/8.11.3; Wed, 4 Apr 2001
  09:00:10 -0400 (EDT)
```

Now it's time to put this into a script that will run from /usr/local/etc/rc.d at boot.

The Script

In order to keep our script up to par with the other scripts put in /usr/local/etc/rc.d, and to keep shutdown from complaining that we're using an "old-style" script, we have to give it start and stop options. An easy way to do this is look look at other scripts in the directory (if there are any others there). 00mysql--client.sh and apache.sh are good examples if you have either of those installed. Here's what the script looks like:

```
#!/bin/sh # define a few variables
sleep=/bin/sleep
su=/usr/bin/su
ssh=/usr/bin/ssh
sshflags="-2 -N -f -L"
host=mail.domain.com
tunnel=9595:mail.domain.com:25
killall=/usr/bin/killall $sleep 4 case "$1"
in
  start)
    $su relay -c "$ssh $sshflags $tunnel
$host && echo -n ' mail relay'
    ;;
  stop)
    $killall -u relay ssh && echo -n ' mail
relay'
    ;;
  *)
    echo "Usage: `basename $0` {start|stop"
>&2
    ;;
esac exit 0
```

This is more elaborate than what's needed, however, I like the fact that changing part of the script can be done by changing the variables at the top. If you don't want to use the variables, you don't have to -- you can just replace them with the commands in the script.

Put the script in /usr/local/etc/rc.d (I've named it relay.sh, make it executable, and it will start your tunnel at boot.

MTA Configuration

The only thing left to do now is set up your MTA to use port 9595 on the local machine as your relay or "smart host". I'm currently using Postfix on my laptop. The relevant part of my /usr/local/etc/postfix/main.cf looks like this:

```
relayhost = [localhost]:9595
```

Add the above line and then restart postfix.

If you're using sendmail, it's a little more involved than that. Here's what you'll need to add to the mc file you use to generate your sendmail.cf:

```
define(`RELAY_MAILER_ARGS', `TCP $h
9595')dnl

define(`SMART_HOST', `relay:localhost')dnl
MODIFY_MAILER_FLAGS(`RELAY', `-k')dnl
```

Now regenerate your sendmail.cf, restart sendmail, and you're ready to use your SSH tunnel to send mail.

Happy mailing,

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

http://www.freebsdzine.org/200104/tunnel.php3

Using Aggregate Functions and Operators in PostgreSQL

By: Branden Williams brw@brw.net

(Part 1 in a multipart series)

Preface:

This series of articles will take into the assumption that the reader can do basic SELECT, INSERT, UPDATE, and DELETE queries to and from a SQL database. If you are not sure on how these functions work, please read a tutorial on how these types of queries work. Specifically if you can use a SELECT query, then you are armed with enough information to read through this document with a high level of understanding. That said, lets get on to aggregate functions!

Summary:

In Part 1 of this series, I will cover how to use the five most common andbasic aggregate functions on PostgreSQL. Those functions are count(),min(), max(), avg(), and sum(). Look for more fun with aggregates andoperators later in this series of articles.

What is an aggregate function?

An aggregate function is a function such as count() or sum() that you canuse to calculate totals. In writing expressions and in programming, youcan use SQL aggregate functions to determine various statistics andvalues. Aggregate functions can greatly reduce the amount of coding that you need to do in order to get information from your database.

```
(Excerpt from the PostgreSQL 7.1 manual)
aggregate_name (expression)
aggregate_name (ALL expression)
aggregate_name (DISTINCT expression)
aggregate_name ( * )
```

where aggregate_name is a previously defined aggregate, and expression is any expression that does not itself contain an aggregate expression. The first form of aggregate expression invokes the aggregate across all input rows for which the given expression yields a non-NULL value. (Actually, it is up to the aggregate function whether to ignore NULLs or not ---

but all the standard ones do.) The second form is the same as the first, since ALL is the default. The third form invokes the aggregate forall distinct non-NULL values of the expression found in the input rows.

The last form invokes the aggregate once for each input row regardless ofNULL or non-NULL values; since no particular input value is specified, itis generally only useful for the count() aggregate function.

Consider this example. You are writing a program which tracks sales ofbooks. You have a table called the "sale" table that contains the booktitle, book price, and date of purchase. You want to know what the totalamount of money that you made by selling books for the month of March2001. Without aggregate functions, you would have to select all the rowswith a date of purchase in March 2001, iterate through them one by one tocalculate the total. Now if you only have 10 rows, this does not make abig difference (and if you only sell 10 books a month you should hopethose are pretty high dollar!). But consider a book store that sells on average 2000 books a month. Now iterating through each row one by onedoes not sound so efficient does it?

With aggregate functions you can simply select the sum() of the book pricecolumn for the month of March 2001. Your query will return one value and you will not have to iterate through them in your code!

The SUM() function.

The sum() function is very useful as described in the above example. Based on our fictitious table, consider the following.

```
table sale (
book_title varchar(200),
book_price real,
date_of_purchase datetime
)
```

Without aggregate functions:
SELECT * FROM sale WHERE date_of_purchase BETWEEN '03/01/2001' AND
'04/01/2001';

This returns all rows which correspond to a sale in the month of March
2001.

With aggregate functions:

```
SELECT SUM(book_price) AS total FROM sale
WHERE date_of_purchase BETWEEN
'03/01/2001' AND '04/01/2001';
```

This returns a single row with a single column called total containing thetotal books sold in the month of March 2001. You can also use mathematical operators within the context of the sum()function to add additional functionality. Say for instance, you wanted toget the value of 20% of your sum of book_price as all of your books have a20% markup built in to the price.

Your aggregate would look like:

```
SELECT SUM(book_price) AS total,
SUM(book_price * .2) AS profit FROM sale
WHERE date_of_purchase BETWEEN '03/01/2001'
AND '04/01/2001';
```

If you look on a grander scale, you will see even more uses for the sum()function. For example calculating commissions, generating detailed reports, and generating running statistical totals. When writing areport, it is much easier to have SQL do the math for you and simplydisplay the results than attempting to iterate through thousands ormillions of records. The count() function.

Yet another useful aggregate function is count(). This function allows you to return the number of rows that match a given criteria. Say forexample you have a database table that contains news items and you want todisplay your current total of news items in the database without selectingthem all and iterating through them one by one. Simply do the following:

```
SELECT COUNT(*) AS myCount FROM news;
```

This will return the total number of news articles in your database.

The MAX() and MIN() functions.

These two functions will simply return the maximum or minimum value in a given column. This may be useful if you want to very quickly know thehighest priced book you sold and the lowest price book you sold (back tothe book store scenario). That query would look like this.

```
SELECT MAX(book_price) AS highestPrice,
MIN(book_price) AS lowestPrice
FROM sale WHERE date_of_purchase BETWEEN
'03/01/2001' AND '04/01/2001';
```

Again, this simply prevents you from selecting EVERYTHING from thedatabase, iterating through each row one by one, and calculating yourresults by hand.

The AVG() function.

This particular aggregate is definitely very useful. Any time you wouldlike to generate an average value for any number of fields, you can usethe avg() aggregate. Without aggregates, you would once again have toiterate through all rows returned, sum up your column and take a count ofthe number of rows, then do your math. In our book store example, say youwould like to calculate the average book price that was sold during March2001. Your query would look like this.

```
SELECT AVG(book_price) AS avg_price FROM sale
WHERE date_of_purchase
BETWEEN '03/01/2001' AND '04/01/2001';
```

Conclusion:

Aggregate functions can greatly simplify and speed up your applications byallowing the SQL server to handle these kinds of calculations. In morecomplex applications they can be used to return customized results frommultiple tables for reporting and other functions. Stay tuned for thenext article in the series!

This article is re-printed with permission. The originals can be found at:
http://www.newbienetwork.net/sections.php?op=viewarticle&artid=25

First Experiences with Red Hat Linux 7.1 on a Workstation

Mike MacCana <mikem@cyber.com.au>

Upgraded my desktop machine from Red Hat 7.0 to 7.1 last week. I thought I might share soem experiences. The machine is an Athlon 900 w/ 128Mb of RAM. Its role in my home network is as a desktop machine, mainly doing 3D accelerated software, X serving, and NFS clienting.

Red Hat still keeps trying to restart X infinitely if you have a bad X config and are using runlevel 5 by default, but give you enough time (just) between retries to try typing something like 'init 3' to go and fix the problem.

The installer now has a laptop options. Asides from that, there's not that much that's different, but it certainly feels better -Red Hat have elininated a lot of the unresponsive elements of their installer - eg, rather than a big pause to transfer the intial system image before the packages are installer, there's a status bar that counts up. This is a small change but it makes ia nice improvement.

The system was seamlessly upgraded, with the only hiccups being Eazel's Nautilus package having some unsatisfied dependencies, and the drivers for my NVIDIA GeForce 2 MX which didn't work with the new kernel (this was to be expected, and fixed easily by downloading 2 packaged drivers from NVIDIAs website).

It was nice to see my Soundblaster live 128 supported out of the box (kernel 2.4 has a wider array of OSS sound drivers than 2.2 did). These Ensoniq 1370 and 1371 chipsets are used throughout Creative's Vibra range of budget cards, so its nifty that they're supported out of the box.

The new Internet config tool is only a 0.4. Its designed to configure ADSL, ISDN, and PPP conections. It works reasonably well, but if it can't find pppd it will simply complain it isn't installed. A prompt to install it right then would have made some sense.

The GUI Firewall editor is nice. This isn't gnome-lokkit (the end user firewall tool from 7.0 that doesn't need ipchains knowledge), its just an elegant interface to ipchains that allows someone to modify rulesets very quickly without having to worry about ipchains syntax.

Using the CD as a rescue disk didn't seem to unmount the filesystems cleanly. This was sort of annoying, because...

- During the second part of the install, I discovered I had a bad second CD (no errors popped up in the installer, but on the console there was read errors)
- I had to end the install then and grab another copy of the CD, then start upgrading my system again
- To resume the upgrade, I needed a clean filesystem. When the installer shut down, it didn't unmount them properly.
- I couldn't use e2fsck from the existing half upgraded copy of Red Hat on my hard disk, because being half upgraded, the system wouldn't boot.
- I couldn't use the Red Hat install CD as a rescue to check my disks, because it too wouldn't unmount the partition uncleanly.
- I ended up needing TomsRTBT to do the e2fsck.

KDE2 continues to be very nice, and the system performs very well. The newer version of update agent also seems to handle dependency issues better, and be mroe stable than, the 7.0 release.

There's no need for kgcc to compile kernels anymore, as kernel 2.4 is more compliant with C standards and thus works well under Red Hat's 2.96 version of GCC.

Even though Red Hat 7.1 is the long kernel 2.4 release, for those who have Red Hat 7.0 and installed all the updates, and kept other packages up to date, its more of a incremental release. But 7.1 is certainly a nicer base to build upon than 7.0, which requires many updates to its basic distribution which are included in 7.1.

Mutt: An e-mail Users Best Friend -- Part One

Steve Manuel <steven.manuel@usa.net>

With all the choices of graphical e-mail clients available to Linux users, the reader may wonder why I would spend time on a terminal based e-mail program. Like other debates on GUI vs. terminal base applications, the answer to this question is "it depends". It depends on what your needs are, what you're use to, and finally, what your preference is. Let me start with why I use Mutt and then why you might like it too.

Like most Linux users, I get quite a bit of e-mail from multiple mailing lists. I used Netscape Messenger for quit a while and, for the most part, was happy with it. However, reading mailing lists became a chore with Messenger because it handled threading poorly ;with no way of changing its behavior. At the same time I needed a way of reading my mail remotely. If I used

Messenger to retrieve my mail at work, I wouldn't be able to read it when I was at home and visa versa. I was familiar with Pine and was going to use it when a fri endsuggested I try Mutt. Since it can be configured to use vi like key bindings, (I love vi), I decided to give it a try.

Mutt can be configured to do about anything with your mail. It handles MIME encoded mail quite well, has PGP support, and can be used as a POP3 mail client. The way it works with mailing lists and threaded mail alone is reason enough for any heavy e-mail user to consider it.

You can configure it to use different signatures depending on who you're sending mail. And you can configure it to view mail in different mailboxes in different ways, including using color highlighting for different parts of an e-mail. If there is some function you want Mutt to have, it can be configured to have that function.

E-mail overview

Before we get going lets review how mail is sent and received on a Linux system and the programs that are involved. Mail is handled by three programs: the MTA (Mail Transfer Agent), the MDA (Mail Delivery agent), and the MUA (Mail User Agent). The MTA is responsible for sending mail (either local or remote) to the recipients' MTA. The MT Ahands off the mail to the MDA. The MDA then delivers the mail to the default spoolfile; usually/var/spool/mail/username. The users MUA then reads the mail in the spoolfile and either leaves it there or moves it to another place for archiving.

Its slightly different for POP3 e-mail accounts. You still use theMTA to send mail, but to receive mail you use a POP3 client. It connects to the POP3 server, sends your username and password, retrieves your mail and delivers it either to your default mail spoolfile or to a folder the POP3 client is configured to use. Mutt can beused as a straight MUA that reads from the mail spool file of yourchoice or as a POP3 client MUA. I will go over how to configure Mutt for either case.

Installation

The first thing you need to do is install Mutt. Luckily, Muttcomes as a standard application on most Linux distributions so chancesare you have it on your system already. To find out, open a terminal window and type mutt -v. If it's installed, you should see something like this:

```
[steve@turin ~]$ mutt -v

Mutt 1.2.4i (2000-07-07)
Copyright (C) 1996-2000 Michael R. Elkins and oth-
ers. Mutt comes with ABSOLUTELY NO WARRANTY; for
details type 'mutt -vv'. Mutt is free software,
and you are welcome to redistribute it under cer-
tain conditions; type 'mutt -vv' for details.
System: Linux 2.4.0-test7 [using ncurses 4.2]
Compile options:
-DOMAIN
-DEBUG
-HOMESPOOL  +USE_SETGID  +USE_DOTLOCK  +USE_FCNTL
-USE_FLOCK
+USE_IMAP  -USE_GSS  -USE_SSL  +USE_POP
+HAVE_REGCOMP  -USE_GNU_REGEX
+HAVE_COLOR  +HAVE_PGP  -BUFFY_SIZE -EXACT_ADDRESS
+ENABLE_NLS  +COMPRESSED
SENDMAIL="/usr/sbin/sendmail"
MAILPATH="/var/spool/mail"
SHAREDIR="/usr/share/mutt"
SYSCONFDIR="/etc"
ISPELL="/usr/bin/ispell"

To contact the developers, please mail to .
To report a bug, please use the muttbug utility.
```

This tells you the version you have (in my case 1.2.4i), the library used for terminal screen control (ncurses 4.2), and then the compile time options. Since my version of Mutt was from an RPM, the output of mutt -v has whatever compile options t hebuilder saw fit to use. I've compiled Mutt from a tarball and its nomore difficult than installing the RPM. If you have any experiencewith compiling from source you will have no problems with Mutt. Ituses the standard ./configure;make;makeinstall routine that you know and love. If you want to know what all the compile options are just type ./configure--help and take a peek. However, the default options arefine for most users.

When installing the RPM, you will need to have the urlview package installed first or you will get a dependency error. If you compile from source, you don't need to have urlview. However, you will not be able to view URLs in your e-mails without it. You will also need either the ncurses or slang package installed as well.

These packages are used by Mutt for screen control. According to the INSTALL file, slang may work better for some people who don't have proper termcap entries. I have used both without any problem. I won't dwell much on installing the package; its very straight forward.

Configuration

The same can't be said for Mutts configuration file: .muttrc. It's this file that gives Mutt its flexibility and configurability; it'salso this file that might give new users problems. The number of options that Mutt has available to it are truly astounding. Mutt allows the user to control about every function Mutt uses to send,receive, and read your mail. As is true with all powerful software, it takes time to understand the features and what they can do for you.

However, once you start using Mutt, you begin to ask yourself, "I wonder if Mutt does...", the answer is almost always yes.

Lets start with the basics of the .muttrc file and along the way Iwill comment on what I feel are good formatting practices so that you can read (and edit) your .muttrc file later. The following is not an exhaustive list of the options Mutt uses (that would take quite a few more articles), just the ones that most users need to get started.

Most of the options are invoked using the "set" or "unset" commands with either boolean or string values, e.g. "set folder = ~/Mail". I put all the "set" commands in one section and the "unset" in another. I make an

exception for some options that are related and work together.

The first options that need to be "set" are the default mail folder, spool file, and e-mail file type:

```
# Mail Folder
set folder = ~/Mail     # Directory that contains
all mail files (mailboxes)
set spoolfile = +Inbox# Default spoolfile
set mbox = +Inbox       # Where mail is appended
to from spoolfile
set mbox_type = mbox    # Type of mail files
set postponed = +Unsent # Where to save postponed
mail
set copy = yes          # Save copies of outgoing
mail?
set record = +Sent      # Where to save copies of
outgoing mail
```

Notice the "+". It's an aliases that means "In the "folder" directory" as set by the "set folder" option. You can use either "+" or "=". The reason you can use both + or = is that "set spoolfile ==Inbox" looks awkward.

We should clarify what "spoolfile" means here; its wherever your MDA saves your mail. Most people have POP3 mail accounts which means that you will need to use some program to retrieve your mail from that POP3 server. Fetchmail is the most popular, but Mutt can also "pop "your mail as well (more on that later). Fetchmail will "pop" or retrieve your mail then send it to the MDA (generally Procmail). If you configured Mutt for POP3, your mail will be retrieved by Mutt th ensaved to the file specified by "set spoolfile".

If you want the mail in your spoolfile moved to another folder after reading then you should "set mbox" to thatfile. I don't need Mutt to move my Mail out of the default spool file because I have Procmail filter my mail and bypass /var/spool/mail/username and deliver it to other folders based on address. Any mail that Procmail can't save to another mail folder goes to ~/Mail/Inbox. Because of this I don't need "setmbox"; just "set spoolfile". "set mbox_type" lets Mutt know what mailbox typeto create for new mailbox folders. Sendmail and Postfix use mbox; Gmail uses Maildir. *Note: If you are converting from Netscape Messenger to Mutt, use mbox. "set copy = yes" isthe default value for this option but I like to explicitly set it for reference. Even though Mutt saves a copy of outgoing mail it doesn't yet know where to save it so, "set record" tells Mutt which file that is.

Now, the personal options:

```
# Personal options
set hostname = "your hostname"
set realname = "your real name"
set signature = ~/.signature
```

These three values are self explanatory. Set them to values appropriate for you. For "set hostname", I set this for the domain name of my POP3 email service. This way the domain name in the "From" header field that the recipient sees is from mye-mail service.

Lets define where our address book (aliases in Mutt speak) islocated:

```
# Aliases (address book)
set alias_file = ~/.mutt-aliases
source ~/.mutt-aliases
set reverse_alias       # Show real name instead
of e-mail address in index
set sort_alias = alias  # Sort aliases by alias
name not email address
```

The only tricky part here is the second line. You have to "source"the file that contains your email aliases otherwise Mutt won't "see"them. I have my aliases sorted by real name as some of the e-mail addresses are not very intuitive. Check "manmuttrc" for other sort options.

```
This section is imperative:
# Prune the headers!
ignore *# Ignore all header info
unignore subject
unignore to
unignore from:
unignore date
unignore CC
hdr_order Date: From: To: CC: Subject:
```

Have you ever looked at all of the header info on an email? Unless you like this sort of info its fairly useless. What most people want to see is who the mail is from, the date sent, and the subject. The first value ignores everything; the rest unignores only the header-info that's important to you. The "hdr_order" allows you to specify the order you see the header info when you read mail. It's one of those cool features that you wouldn't think are important until you use it.

Tell Mutt what folders to look in for new mail:

```
# Mailboxes
mailboxes =Inbox
mailboxes =Mailing-Lists/Kickstart-list
mailboxes =Mailing-Lists/LinuxWorld
mailboxes =Mailing-Lists/PPA-Devel
mailboxes =Mailing-Lists/PPA-Users
mailboxes =Mailing-Lists/RPM-list
```

The "mailboxes" variable will tell Mutt what mailboxes can receive mail and where Mutt should check for new messages. You can cycle through these mailboxes by pressing the space bar after changing folders. This allows you to keep an eye on certain mailboxes for new mail. *Note: Mutt will not put mail in these folders; you will need something like Procmail to do that.

If you want to use Mutt as a POP3 mail client...:

```
set pop_host =[your.pop3.hostname]
set pop_user =[your pop3 username]
set pop_pass =[your pop3 password]
set pop_delete =yes# Save mail on server or
not
```

If you don't want to use a POP3 mail retriever like Fetchmail, you can configure Mutt as a POP3 mail client and have it retrieve your mail. The options are straight forward. The last one will depend on whether or not you want to save your mail on the POP3 server. Since most POP3 servers have a limit on how much mail they will hold for you, I suggest you have Mutt delete mail after its retrieved it. It doesn't take long for that limit to be reached.

With these options set in your .muttrc file you can use Mutt. There are many, many more options and the curious reader should consult "man 5 muttrc" for the rest. Now that wehave the .muttrc file created you will

have to create your alias file.A simple "touch [nameofaliasfile]" is all that isneeded.

In my next article I will go over the functions of Mutt and how to navigate your way around your mail. Again, if you want to know more ,the man pages for muttrc and mutt are quite helpful. I also recommend the Mutt manual. On RedHat systems it's installed in "/usr/doc/mutt-1.2.x". This document has all the muttrc options and plenty of instruction on how to use Mutt.

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

http://www.linuxnovice.org/main_software.php3?VIEW=VIEW&t_id=146

---

ssh suite: Sftp, scp and ssh-agent

Author: Matteo DellOmodarme <matt@martine2.difi.unipi.it>

The aim of this article is to provide an introduction to some useful programs in the SSH suite, i.e. sftp, scp, ssh-agent and ssh-add.

In the following we suppose that sshd2 daemon is well configured and running.

Sftp and scp overview

Let's focus our attention on sftp and scp.

The first one (Secure File Transfer) is a ftp-like client that can be used in file transfer over the network.

It does not use the FTP daemon (ftpd or wu-ftpd) for connections, allowing a significant improvement in the system security. In fact, monitoring some logs file of our systems, we noted that about 80% of attacks in last month was against ftpd daemon. The use of sftp prevents all these tries since it permits to stop the potentially dangerous wu-ftpd.

The second (Secure Copy) is used to copy files over the network securely. It is a replacement for rcp insecure command. Sftp and scp do not require any dedicated daemon since the two programs connect to sshd servers.

In order to use sftp and scp you have to insert the following line in the configuration file `/etc/ssh2/sshd2_config:`

```
subsystem-sftp                  sftp-server
```

after this modification you must restart sshd.

So you could use sftp and scp only to connect to hosts where sshd is running.

Sftp

Sftp uses ssh2 in data connections, so the file transport is as secure as possible. There are two main advantages in using sftp instead of ftp:

1. Password are never transferred in clear text, preventing any sniffer attack.

2. Data are encrypted during the transfer, making difficult to spy or modify the connection.

The use of sftp2 is really simple. Let's suppose that you would connect via sftp to your account myname on host1. In order to do that use the command:

```
sftp myname@host1
```

some options could be specified from the command line (see the sftp manual page for a complete report).

When the sftp2 is ready to accept commands, it will display a prompt sftp>.

In the sftp manual page there are a complete list of the commands which the user can use; among them there are:

quit:
Quits from the application.

cd directory:
Changes the current remote working directory.

lcd directory:
Changes the current local working directory.

ls [-R] [-l] [file ...]:
Lists the names of the files on the remote server. For directories, the contents of the directory are listed. When the -R option is specified, the directory trees are listed recursively. (By default, the subdirectories of the argument directories are not visited). When the -l option is specified, permissions, owners, sizes and modification times are also shown. When no arguments are given, it is assumed that the contents of . are being listed. Currently the options -R and -l are mutually incompatible.

lls [-R] [-l] [file ...]:
Same as ls, but operates on the local files.

get [file ...]:
Transfers the specified files from the remote end to the local end. Directories are recursively copied with their contents.

put [file ...]:
Transfers the specified files from the local end to the remote end. Directories are recursively copied with their contents.

mkdir dir (rmdir dir):
Tries to create (destroy) the directory specified in dir.

sftp2 supports glob patterns (wildcards) given to commands ls, lls, get, and put. The format is described in the man page sshregex.

Since sftp use encryption there is drawback: the connection is slower (about a factor of 2-3 to my experience), but this point is of marginal interest considering the great security benefits.

In a test conduced on our local network a Network Sniffer was able to catch a mean of 4 password by hour, from ftp connections. The introduction of sftp as standard protocol for transfer file across the network could eliminate this security problem.

Scp

Scp2 (Secure Copy) is used to copy files over the network securely. It uses ssh2 for data transfer: it uses the same authentication and provides the same security as ssh2. It is probably the simplest way to copy a file into a remote machine. Let's suppose you want to copy the file filename contained in the directory local_dir to your account myname on the directory remote_dir on host host1. Using scp you could enter from the command line:

```
scp local_dir/filename myname@host1:remote_dir
```

In such a way the file filename is copied with the same name. Wildcards can be used (read more about those from sshregex man page).

The command:

```
scp local_dir/* myname@host1:remote_dir
```

copies all files from directory local_dir into the directory remote_dir of host1.

The command:

```
scp myname@host1:remote_dir/filename .
```

copies the file filename from remote_dir on host1 to the local directory.

Scp supports many options and allows copies between two remote systems as in the following example:

```
scp myname@host1:remote_dir/filename  \
myname@host2:another_dir
```

See its manual page for a complete presentation.

Obviously, using scp, you must know the exact directory tree of the remote machine, so in practice sftp is often preferred.

ssh key management

SSH suite contains two programs useful to manage authentications keys, allowing the user to connect to a remote system without specifying a password or even a passphrase. These programs are ssh-agent and ssh-add.

ssh-agent

From the manual page of ssh-agent we can read:

"ssh-agent2 is a program to hold authentication private keys. The idea is that ssh-agent2 is started in the beginning of an X-session or a login session, and all other windows or programs are started as children of the ssh-agent2 program (the command normally starts X or is the user shell). The programs started under the agent inherit a connection to the agent, and the agent is automatically used for public key authentication when logging to other machines using ssh".

There are two way to use ssh-agent depending on that you are running xdm or not.

In the first case you should edit .xsession file, placed in the $HOME directory. There are two possible procedures:

Copy .xsession to .xsession-stuff and modify .xession in such a way it contains only the line:

```
exec ssh-agent ./.xsession-stuff
```

Alternatively you could edit .xsession file and search for each line containing the expression "exec program". Modify
these lines to the form "exec ssh-agent program".

Log out from your X-session and restart it. ssh-agent will start the X-session as its own children and wait for ssh key to insert in its database.

If xdm is not running the procedure to use ssh-agent is simpler because you can start your X session using the command:

```
ssh-agent startx
```

In such a way you have ssh-agent properly running.

ssh-add

Once ssh-agent is correctly in place you could add identities in its database using the command ssh-add. You could add identities only from processes which are children of a ssh-agent ancestor otherwise the following error message is displayed:

Failed to connect to authentication agent - agent not running?

The use of ssh-add is simple: from the command line issue the command:

```
ssh-add
```

ssh-add scans the file $HOME/.ssh2/identification
which contains names of the private keys that are to be used in authentication. If this file doesn't exist, the standard name for the private key is assumed (i.e. $HOME/.ssh2/id_dsa_1024_a).

If any public key file requires a passphrase, ssh-add asks for the passphrase from the user as in the following example:

Adding identity:
`/home/matt/.ssh2/id_dsa_1024_a.pub`

Need passphrase for
/home/matt/.ssh2/id_dsa_1024_a (..)

Enter passphrase:

You could obtain a list of all identities currently represented by the agent using the command

```
ssh-add -l:
```

Listing identities. The authorization agent has one key:
```
id_dsa_1024_a: 1024-bit dsa, (...)
```

Conclusions and useful links

Many users of telnet, rlogin, ftp might not realize that their password is transmitted across the net unencrypted, but it is.

The use of some secure protocols could allow a secure transmission over an insecure network. SSH, encrypting all traffic, effectively eliminates eavesdropping, connection hijacking, and other network attacks.

These articles are only an introduction to the SSH suite; more about this topic could be found in the manual pages of ssh,
sshd and sftp.

You could get SSH suite from:
www.ssh.com/products/ssh/, SSH master site or
from a mirror site.

Here you could also find some very interesting information about SSH technology and cryptography in general in the Tech corner.

Otherwise you could check www.openssh.com where you could download openssh implementation of SSH protocol. The portable version is at
www.openssh.com/portable.html.

You could also read the openssh FAQ: www.openssh.com/faq.html.

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

http://www.linuxgazette.com/issue64/dellomodarme.-html

Consultant's Perspective: Mandrake Single Network Firewall

Ng Kai Hoe Raymond <raymond@aeonxe.com>

Yesterday was an eventful day for me. As a security consultant, my services have been retained by my client to setup a firewall to protect my client's office network against potential external hackers. Guess what? I installed Mandrake's Single Network Firewall for my client. That is a distribution released by Mandrake just weeks ago.

I shall talk about Single Network Firewall from a consultant's viewpoint. That is, I will talk about how its function helps me as a security consultant to accomplish my work. I will talk little on its technical merits, as I personally don't think wrestling with Single Network Firewall for just one day qualifies me to talk about its technical flaws, or merits, or the lack thereof.
I shall talk about the following items

a) Installation
b) Configuration
c) Funny quirks here and there
d) Price and Time Efficiency
e) Final Evaluation

Installation

Installation is a breeze with Mandrake Single Network Firewall, just that it has some funny assumption which I am not ready to live with yet. Like for example, it is using eth0 as the LAN interface of the Single Network Firewall - it has delayed my complete installation and configuration by a sordid one hour. Functionally, traditionally and culturally, I have always installed eth0 as the WAN interface, and the rest as the LAN (or DMZ) interfaces, and that was the root of my installation sorrows.

Let me explain, Single Network Firewall's web based configuration only allows a browser lying on a machine on the LAN interface to administer the firewall. Any requests coming from the WAN interface will be completely denied. As a matter of fact, any machine on the WAN interface cannot even ping the firewall. What I did was to connect eth0 to the WAN interface and connect the LAN interface to eth1, when it should be the other way round, and tried as I had, I could not configure the firewall from a machine lying on the eth1 interface, because Single Network Firewall designated it to be the WAN interface. When that happened to me yesterday, I thought that there was something wrong with other parts of the network, only to find out later that the firewall probably has denied all ping requests to their WAN interface. Quite some time was wasted there.

Other than that, I am pretty satisfied with the intuitive setup that Mandrake's Single Network Firewall presents to their users. It has a small footprint (200 over MB only), and I am already thinking of installing Single Network Firewall on a solid state hardisk. That is the only way to ensure a very good uptime, because the solid state hardisk has much higher MTBF than a hardisk.

Configuration

Configuration of a firewall is easy for Single Network Firewall. It is done through a web interface with the URL taking the form of https://ip:8443 where ip is the IP address of your firewall. As you can see, the web

interface is protected by strong SSL-cryptography and you can configure it to be accessed by only requests coming from a particular NIC card - that functionally means that you can deny configuration attempts from the Internet altogether, and nobody can make sense of your configuration sessions.

For those who are interested in how the configuration looks like, please refer to this link. I shall not talk much of that here.

Configuration obviously also includes the technical realisation of all the security policy that I have discussed with my client beforehand. And their web interface is superb in the sense that it simplifies my job for me immensely, I just need to specify in high level ideals on how I want the firewalling to be. Like for example, "I want to deny all smtp request from LAN to WAN", and there are intuitive checkboxes that I can click on to immediately realise technically these policies. I do not even remember inputing any ipchains, iptables or ipfwadm commands. To put it plainly, I do not need to know these in the first place.

There is also an option to configure Intrusion Detection Systems, both Snort and Prelude. No ugly compilation and no downloading of rulesets from Snort site. Just turn on the Intrusion Detection Systems, and you have a firewall which logs down all possible hack attempts.

All in all, configuration screens are intuitive and there are not too many problems moving around.

Funny Quirks Here and There

There are still a few quirks here and there that I am still not very used to. Since there are not too many logical relationship between the quirks, I shall list them in point forms.

1) I am supposed to have a DMZ on a third interface (eth2), but this Single Network Firewall seems to only have the idea of LAN and WAN. Both the DMZ (eth2) and the office network (eth1) are treated as LAN. So if you want to deny all SMTP requests from eth1 but allow all SMTP requests from eth2, you are in trouble. There are no adminstrative functions that address the kind of granularity that you need. I may need to know something about ipchains in the end.

2) After the first stage setup, a machine from the DMZ (eth2) cannot even ping a machine on the office network (eth1). That is a nice thing to have, but there is not any administrative function in the configuration tool that allows me to allow traffic to and fro from eth1 and eth2.

With that I end my case on the quirks. It seems bizzare enough, I may have to revise my ipchains to actually complete this assignment.

Price and Time Efficiency

Price efficiency is there definitely. You are using an Open Source product, and you only need to pay for the consultant's manhours. Not to mention you get a ready made appliance with all the IDS installed, and a friendly administration tool that anyone can take over after the consultant has installed the Single Network Firewall. A similar commercial product would cost you 20 k USD, perhaps.

The other part of the price probably comes from hardware, which is really cheap nowadays.

Time efficiency for the security consultants is good. You spend so little time installing the distro, you find out that the things that require you to compile are not longer needed - I do not need to install Snort or Prelude, it comes with Single Network Firewall out of the box.

Final Evaluation

Let me just do a evaluation based on these criterias. 0 stars is lousiest, while 5 star is excellent.

On ease of installation - 5 stars

On ease of configuration - 4 stars
Beauty of the web administrative interface - 4 stars

On diversity of administrative functions - 2 stars

Range of software (IDS, WebProxy, etc) - 5 stars

System Monitoring - 3 stars

With that I thank all my readers. And if there is a factual error which I have made above, I am open to corrections. Please send me the errata if you can spot any errors.

Thanks and Best Regards

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

http://www.aeonxe.com/article.php?story=20010704-040418557

Summary of Minutes from AUUG Exec Meeting

By: Liz Carroll <busmgr@auug.org.au>

10:00am - 4:00pm, 3rd February, 2001
UTS, Sydney

Attendees:

Sarah Bolderoff	SB
Greg Lehey	GL
Luigi Cantoni	LC
David Newall	DN
Elizabeth Carroll	EC
Michael Paddon	MP
Alan Cowie	AC
David Purdue	DP
Peter Gray	PG

Notetaker:

Elizabeth Carroll	EC

1. Apologies

Malcolm Caldwell

2. President's Report

Having come through our quiet period, I have little to report.

In November we ran two successful events - AOSS and the Security Symposium. Both events were profitable.

We will need to look at AOSS III - if we want to bring this back to a schedule, we should be aiming for May/June.

Which brings us to the annual conference. I think we all know how important it is to us that this be successful - but we face some real challenges, not least the lack of a programme chair.

3. Secretary's Report

Interesting the percentage changes - membership growing gently - two peaks being SA and TAS, reflection from activities in those areas. Membership is stable - gut feel that we have approximately 500 core members - should be aiming for approx 800 to achieve some of the goals we have been aiming for - this however is MP's subjective opinion.

4. Treasurer's Report

Looking good - we are not eating into our funds, even though we are not holding a conference (during tax year 1 July 2000 - 30 June 2001). All events have been positive. This shows we should be running more of these events and getting them out to our members. Nothing is running over budget. Income has been over budget - nothing much more we can do on saving money. Organisation is financially sound. No problem to finance the conference.

5. Business Manager's Report

There are still 3 outstanding accounts from AUUG2K - discussion ensued on how to deal with late payments in future. We shall introduce a surcharge for processing purchase orders - at Business Manager's discretion we may refuse purchase orders. At 180 days, we should start legal action.

Events run recently were successful.

6. Minutes of Previous Meeting

MP spent a week in December following up on the Company returns - there has been a break in activity - will be following up on it.

7. Action Items

Action Items were addressed.

8. Events

Motion of thanks for the organizers of the Security Symposium

A draft timetable for future events was discussed.

AUUG 2001

There was discussion regarding the venue, possible speakers, sponsors, Programme Committee and general issues relating to the conference

9. Web Server

DN and DP have spoken with Simon Hackett - he is ready to start set-up with the server.

DN and DP have spoken with Simon Hackett - he is ready to start set-up with the server.

Chapters should have the choice of using this web server or their own.

Action: DP will notify the chapters of our current activities, namely Canberra.

10. AUUG Use of UNIX

Discussion ensued on the use of UNIX by AUUG, following details provided by GL.

11. Exec Communication / Decision Making Process

Discussion ensued on the use of Exec Member's using email. We should possibly look at issues that require a decision being made on, to be sent to DP, for him to put it to the members for a vote, in order to effect a resolution on issues.

There will be 48 hours between an AUUGFINAL: and AUUGVOTE: There will be a maximum of 7 days for the AUUGVOTE:

The subject for a resolution from David will start: AUUGVOTE: This will be preceded by an email called: AUUGFINAL:

12. AUUG Involvement with SAGE-AU - Conference

DP stated that for many years we have offered SAGE a stream at AUUG Conferences. Last year was the exception - due to their conference being too close to AUUG's. AUUG should now commence discussions with SAGE on this topic.

GL stated that he met with Andrew van der Stock in Adelaide. Following conversation, as AUUG and SAGE have approx 20 members each in SA, it was suggested joint meetings be held. Subsequently, their annual conference is being held in Adelaide in June - GL offered some AUUG help.

13. Special Interest Groups

BSD SIG has not happened yet, MP is currently working on this.

MP feels that SIGs are the way forward. Most of our members and potential members are mainly located in Sydney and Melbourne. In both these locations, 'Chapters' do not work - this does not help raise our membership numbers in these
areas.

SIG's that could work include:

- BSD
- Security
- Network

When we say SIG, it should include:

- 1 Event per year
- Various activities
- Mailing list

Ownership of SIG's to date.

- BSD - MP
- Security - AC

14. IT Council for SA

GL became aware of them just prior to AOSS2. Basically, they are a non-profit group representing the IT industry in SA. They give help to the government regarding IT policy. There is a policy for membership, however this is not enforced. There is an application form to be completed - AC suggests that a couple of Exec Committee members should join. However there are no individual members - just corporates. Therefore, we could look at 'AUUG' becoming a member - and if there are meetings we can send an Exec member.

15. Business Manager to visit USENIX

GL stated that EC is currently working on her own without much feedback from others. GL feels that EC could learn from helping out at USENIX, in the central office.

16. Advertising and Press Releases

GL suggested a Press Reception to let the media know what AUUG is currently doing.

Discussion ensued on various ideas.

MP suggests we give the media 'copy' which they can use. This is designed to be immediately useable to fill space.

DP suggests that in July / August we have a conference launch - invite the press, by that stage we will know our speakers so we can this to speakers. MP suggests we send them a pack.

17. AUUGN

Current issue in progress - all going well.

18. Chapters

There was discussion on the viability of a SA chapter. DN spoke of an idea of having a 'beer and pizza' night every month on a trial basis for maybe 6 months.

Motion that AUUG allocate $50 a month for the next 3 months to be administered by SB to be used to sponsor chapter establishment activities in SA. Moved DN, Seconded MP. Carried

19. Membership and Renewals

Previously covered under Secretary's Report.

20. Other Business

Action MP and DN: Next edition of AUUGN must carry Election Forms - they will also be upon the web.

National Installfest - means in a number of major centres we would have an Installfest - on the same day. Suggestion has been the 18 or 25 August.

Action - to be discussed over email.

Meeting Closed: 4.25pm
Next Meeting: 5 May 2001, Sydney

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 or call the AUUG office on 1-800--625655.

Membership Application

FRONT

Membership Application

BACK

Volume 22 • Number 2
July 2001

Features:

News:

Regulars: