

# AUUGN

The Journal of AUUG Inc.

Volume 24 • Number 4  
December 2003

## Features:

Collaborative Virtual Workspaces and their advantages	8
A short Guide to Installation of CVW Under Linux	10
Building Nefarious 3/5	13
2003 and Beyond, Pt. 2	17
GIMP: Make Your Own Brushes	29
Linux Radio Station	33
Working with ACLs in FreeBSD 5.x	36
Mandrake 9.1 - Raising the Bar in GNU/Linux	40
Distribution Usability	47
CyberInsecurity: The Cost of Monopoly (part 1 of 2)	50
An Overview of Sentinel3G	53
Supercluster Redux	55
The Linux Desktop Underrated	55

## News:

Public Notices	5
AUUG Conference 2004: Call for Papers	6
linux.conf.au 2004	7
AUUG: Corporate Members	35
AUUG: Security Symposium 2004: Call for Papers	36
AUUG: Chapter Meetings and Contact Details	

## Regulars:

President's Column	3
My Home Network	4
AUUGN CD-Rs in this issue	5





#### AUUG Membership and General Correspondence

##### The AUUG Secretary

AUUG Inc  
PO Box 7071  
Baulkham Hills BC NSW 2153  
Telephone: 02 8824 9511  
or 1800 625 655 (Toll-Free)  
Facsimile: 02 8824 9522  
Email: [auug@auug.org.au](mailto:auug@auug.org.au)

##### AUUG Management Committee

Email: [auugexec@auug.org.au](mailto:auugexec@auug.org.au)

##### President

Greg Lehey  
PO Box 460  
Echunga, SA, 5153  
Bus. Tel (08) 8388 8286, Mobile 0418 838 708, Fax (08) 8388 8725  
<[Greg.Lehy@auug.org.au](mailto:Greg.Lehy@auug.org.au)>

##### Vice-president

David Purdue  
Sun Microsystems  
Level 6, 476 St Kilda Road  
Melbourne, Victoria, 3004  
Phone: +61 3 9869 6412, Fax: +61 3 9869 6288  
<[David.Purdue@auug.org.au](mailto:David.Purdue@auug.org.au)>

##### Secretary

Adrian Close  
ActiveSky Pty Ltd  
Level 1, Earle Plaza Cnr. Price & White Streets  
Nerang QLD 4211  
Business Telephone: +61 7 5596 2277  
Business Fax: +61 7 5596 6099  
Mobile: +61 417 346 094 <[adrian@auug.org.au](mailto:adrian@auug.org.au)>

##### Treasurer

Gordon Hubbard  
Custom Technology Australia Pty Ltd  
Level 22, 259 George Street, Sydney NSW 2000  
Bus Tel: 02 9659 9590, Bus Fax: 02 9659 9510  
<[Gordon.Hubbard@auug.org.au](mailto:Gordon.Hubbard@auug.org.au)>

Jonathon Coombes  
Cybersite Consulting Pty Ltd  
34 Newcastle Road, Wallsend NSW 2287  
Business Telephone: +61 2 4965 6989  
<[Jonathon.Coombes@auug.org.au](mailto:Jonathon.Coombes@auug.org.au)>

Andrew Frederick Cowie  
Operational Dynamics  
GPO Box 4339  
Sydney, NSW, 2001  
Telephone: +61-2-9977-6866  
<[Andrew.Cowie@auug.org.au](mailto:Andrew.Cowie@auug.org.au)>

Steve Landers  
Digital Smarties  
PO Box 717 Willetton WA 6155  
Business phone: +61 8 9313 6868  
Business fax: +61 8 9313 6077  
<[Steve.Landers@auug.org.au](mailto:Steve.Landers@auug.org.au)>

Stephen Rothwell  
IBM Australia Linux Technology Centre  
8 Brisbane Ave Barton ACT 2600  
Business Telephone: +61 2 62121169  
<[Stephen.Rothwell@auug.org.au](mailto:Stephen.Rothwell@auug.org.au)>

Michael Still  
12 Bushby Close Gordon ACT 2906  
Phone: +61 414 382 568 <[mikal@auug.org.au](mailto:mikal@auug.org.au)>

##### AUUG Business Manager

Elizabeth Carroll  
AUUG Inc  
PO Box 7071  
Baulkham Hills BC NSW 2153  
<[busmgr@auug.org.au](mailto:busmgr@auug.org.au)>

## Editorial

Con Zymaris [auugn@auug.org.au](mailto:auugn@auug.org.au)

*Overheard, on the cobbled road back to Athens from Piraeus, the following...*

GLAUCON: So, tell us Socrates, of this curious Microsoft-SCO funding paradox.

*Socrates, grey-beard flowing, lifts his tunic as he steps over a pot-hole, frowns and responds.*

SOCRATES: It is a very curious situation indeed. Here we have one of the great avowed opponents of Unix, happily handing over many millions in exchange for property access rights to Unix code and technology that they do not need nor are they likely to ever use. We must ask ourselves a question; what does a firm like Microsoft gain from licencing this technology?

ADEIMANTUS: But surely Socrates, Microsoft wants to licence SCO's technologies to include in their product base, such as their Unix Services for Windows.

SOCRATES: If that were indeed true, my dear friend, if they really did need to licence this technology as part of a shipping product, would they not have taken the licencing step *before* they produced, released and sold that product? No, the answer lies much deeper.

GLAUCON: Perhaps they want to include the SCO technology in that shipping product now? Did not Homer make mention of SCO as the final resting place of all the ancient technologies liberated by Prometheus from the Titans? Perhaps SCO technology is still worth using?

SOCRATES: Firstly, there appears to be no SCO technology in Microsoft's product, only their own and GNU technologies. I ask you, my friends; who amongst you would be the first to replace GNU with SCO? You will recall that it was Homer himself who recounted the story of how the Berkeleyites created a great wooden horse with which they cleverly out-manoeuvred SCO's predecessors in that famous ancient battle, marking the beginning of the demise of the codebase SCO inherited. No, my friends. there is one singular and very deliberate reason why Microsoft has chosen this time to provide revenue to SCO and it has nothing to do with code or technology licencing.

ADEIMANTUS: Perhaps, Socrates there are actions and decisions that our eyes and ears are not privy to? Perhaps this area is not worth our examination?

SOCRATES: The unexamined industry is one that is not worth working in. No, we will ask and we will seek responses to our questions. What does Microsoft gain by supplying SCO with money presently? Is there a target for their actions that they are using SCO to attack, by proxy, as it where? Why undertake this attack by stealth? Are there ethical and moral implications to these hidden actions? Answer this, my friends and we will better understand this industry.

# Contribution Deadlines for AUUGN in 2004


Volume 25 • Number 1 – March 2004: **February 15<sup>th</sup>, 2004**

Volume 25 • Number 2 – June 2004: **May 15<sup>th</sup>, 2004**

Volume 25 • Number 3 – March 2004: **August 15<sup>th</sup>, 2004**

Volume 25 • Number 4 – December 2004: **November 15<sup>th</sup>, 2004**

AUUG Incorporated gratefully acknowledges the support of its corporate sponsor:



---

## AUUGN Editorial Committee

The AUUGN Editorial Committee can be reached by sending email to: [auugn@auug.org.au](mailto:auugn@auug.org.au)

Or to the following address:  
AUUG Inc  
PO Box 7071  
Baulkham Hills BC NSW 2153

Editor:  
Con Zymaris

Sub-Editors:  
Frank Crawford, Mark White

Contributors:  
This issue would not have happened without the transcription and editorial efforts of Gary R. Schmidt\* <[grschmidt@acm.org](mailto:grschmidt@acm.org)>, Rik Harris <[rik@kawaja.net](mailto:rik@kawaja.net)>, Raymond Smith <[zzrasmit@uqconnect.net](mailto:zzrasmit@uqconnect.net)>, David Lloyd <[lloyd0076@adam.com.au](mailto:lloyd0076@adam.com.au)>, Peter Sandilands <[peter@sandilands.vu](mailto:peter@sandilands.vu)>, Grahame Bowland <[grahame@ucs.uwa.edu.au](mailto:grahame@ucs.uwa.edu.au)>, Cameron Strom <[c.strom@statscout.com](mailto:c.strom@statscout.com)>, Steve Jenkin <[sjenkin@canb.auug.org.au](mailto:sjenkin@canb.auug.org.au)>, Andre Joannis <[andrej@marpware.com](mailto:andrej@marpware.com)>, Miles Goodhew <[mgoodhew@internode.on.net](mailto:mgoodhew@internode.on.net)>, John Chrisoulakis <[John.Chrisoulakis@aad.gov.au](mailto:John.Chrisoulakis@aad.gov.au)>, Daniel O'Connor <[doconnor@gsoft.com.au](mailto:doconnor@gsoft.com.au)>, Colin Charles <[byte@aeon.com.my](mailto:byte@aeon.com.my)>

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:

<http://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  
PO Box 7071  
Baulkham Hills BC NSW 2153

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

---

## Mailing Lists

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

### Disclaimer

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

### Copyright Information

Copyright © 2003 AUUG Inc.

All rights reserved. Portions © by their respective authors, and released under specified licences.

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

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

---

# President's Column

Greg Lehey <[Greg.Lehey@auug.org.au](mailto:Greg.Lehey@auug.org.au)>

## DEATH OF UNIX, FILM AT 10

One of the most astounding events in the UNIX world this year has been SCO's lawsuit against IBM, Linux users, maybe BSD users, and the industry at large. I've spoken at length on the subject, and I don't intend to repeat myself here. But it's worth noting that what we're seeing here is proprietary UNIX in its death throes.

UNIX is now 34 years old, two thirds of the history of commercial (non-research) computing. Things have changed a lot in that time, and they're still changing at a frightening rate. In particular, though, the use of computers has risen from the activities of a privileged and clandestine few to a way of life not only for us in the industry, but also for the man in the street. It continues to be big business.

In the early days, software was proprietary--so proprietary that there was no need to protect the source code. Apart from the obvious reason that it was easier for vendors to develop software when they had complete control, it helped lock customers to the vendor. But even when UNIX started off, some trends to standardization had become apparent. Operating systems were still platform specific, but many programming languages were not, and the tendency was for general purpose programming languages to become more portable between systems.

This still hasn't happened completely in the operating system world. The only operating system used to a significant extent across vendor platforms is UNIX (sometimes pronounced "Linux"). This makes the user's life easier and vendor lock-in harder, just like portable programming languages did thirty years ago.

There are differences, though: thirty years ago there was a rule of thumb that writing a compiler took ten man years of work. Productivity has improved since then, offset to some extent by the resultant ability to increase the complexity of the compilers (and, in some case, the languages), and it's now possible to write a compiler a lot faster than that. It would be as good as impossible, though, to write a general purpose operating system in anything like that time. Linus Torvalds and friends had the best chance: they started with a relatively known design, little experience and lots of enthusiasm, and they were eminently successful. But the effort they expended must go into the thousands of man years.

That's all the more reason to share the results. The history of UNIX is full of stories where the technical people wanted to share their code and the lawyers stopped them, or at least tried to do so. The result is a number of UNIX flavours derived (a long time ago) from UNIX System V. Good operating systems, all a little different, but usable. They're also a real maintenance

liability.

The other side of the coin is BSD UNIX, now free (we think). Ten years ago, it was clear that it had no future: the future was reserved for System V. It's amazing how differently a sure thing can turn out to be.

The real winner, though, is undoubtedly Linux. That's nothing new. The thing that's interesting to note now is that it has effectively subsumed proprietary UNIX. Yes, the owners of UNIX have come out fighting, but I don't think any of us have any doubt how it is going to finish. My bet is that if UNIX System V survives at all, it will be because it is released as Open Source. More likely, though, this won't happen until it is completely irrelevant. Two years ago Caldera released "Ancient UNIX" as Open Source, when the newest part of the code base was about 20 years old and completely irrelevant. It's possible that something similar might happen with System V and UnixWare. That would be a pity: if it were to happen soon enough, it might even give SCO a chance of survival.



# Public Notices

---

Upcoming Conferences & Events

## Linux and Open Source in Government

The challenges

12-13 January 2004

Adelaide,

<http://lca2004.linux.org.au/ocgconf/>

## linux.conf.au 2004

Jan 14 until Jan 17, 2004

Adelaide

<http://lca2004.linux.org.au/>

## AUUG Security Symposium 2004

20 February 2004,

Canberra

<http://www.auug.org.au/events/2004/security/cfp.html>

# My Home Network (December 2003)

---

By: Frank Crawford <[frank@crawford.emu.id.au](mailto:frank@crawford.emu.id.au)>

What a difference a few months can make, and hasn't the Open Systems World changed since my last column. In that time we have had, in no particular order, SCO suing everyone having anything to do with Linux (and threatening BSD as well), Red Hat moving to a far more commercial stance, SuSE being bought by Novell and a number of Open Source distribution projects having their servers broken into.

Of course all this makes for an exciting world, but also gives a lot of opportunities for fun. Now, I'm sure that elsewhere in AUUGN there will be plenty of discussion on SCO's claims, so I skip it, but I will follow up on the changes to Red Hat's policy.

Now for those who haven't heard, Red Hat have declared an end-of-life and end-of-support for all free Red Hat Linux software, i.e. Red Hat Linux 7.X and 8.0 from 31st December 2003 and Red Hat Linux 9 from April 30, 2004. In addition Red Hat released Red Hat Enterprise Linux Version 3 (RHEL3), which they will continue to support for the future. Unfortunately, RHEL3 costs a reasonable (or unreasonable, depending on your point of view) amount, as do the ongoing support charges. Of course their support of Red Hat Linux has been the thing that has given Red Hat its position within the Open Source Community.

Fortunately, all is not lost, especially, as with most Open Systems, if you don't like something they do, go and find a version you do, since there are usually half a dozen around. Now in terms of Linux distributions, there are a number of major ones to download, as well as many minor specialist ones. For example you can choose from SuSE, if you wish commercial support, as well as Debian, if you want broad user support. (Of course, Novell's purchase of SuSE may well change all this.) Now if you want something a bit less commercial you could try Slackware, and you can even try Mandrake.

If there is one drawback with Open Source and public domain software, it is the lack of documentation and information, and the change in Red Hat's direction is no different. While Red Hat's publicity machine have been relatively quick to announce all their changes, there is another group that has been a bit slower. While Red Hat have put most of their support behind their commercial product, they are also supporting The Fedora Project, an "Open Source community" project "for developers and early high-tech enthusiasts using Linux in non-critical computing environments" (Red Hat's definition).

In fact the Fedora Project is really the successor to Red Hat Linux, but rather than being support as a commercial entity by Red Hat it is now being supported by the Linux community at large. Red Hat are providing some supporting infrastructure and

## ALLOCPSA

Online Professional Services Automation



allocPSA is a suite of integrated applications (a Professional Services Automation suite) designed for services-based organisations. It enables services personnel to become more productive and profitable by improve their efficiency through increased utilisation and productive time, better planning and through integrated knowledge management.

allocPSA is a complete organisation-wide business solution that attempts to integrate all practice groups and functions in a professional services company into a single computer system. allocPSA consists of numerous software modules for business areas such as resource planning, project management, time and expenses, integration with existing invoicing, book-keeping and payroll systems, faults, messages, announcements, reminders and knowledge management, collaboration, services supply chain, human resources and staff skilling and management of cost-centres.

allocPSA offers total integration between the modules, as well as an open architecture and integration with your existing processes and software. allocPSA is an online, web-based suite, and as such, can be deployed to dozens of users with zero client-side installation, within your Local Area Network, Wide Area Network or as an extranet application to authorised personnel and partners from anywhere on the Internet

allocPSA is supplied as a self-contained server appliance, which contains the operating system platform, SQL server, web server and core allocPSA applications.

allocPSA has been designed to offer your firm the ultimate in flexible deployment integration and ongoing control. To that end, the complete source code for allocPSA is supplied under the GPL Open Source licence. You are then able to extend, modify or maintain this code, should you so wish, or hire 3<sup>rd</sup> parties (Cybersource, or others) to undertake this work for you.

The fee schedule for allocPSA includes a pre-purchased quanta of installation, customisation and support services. A fully-functioning live demo is available upon request.

Web: <http://www.cyber.com.au/cyber/product/allocPSA/>  
Phone: +61 3 9621 2377 Mail: [info@cyber.com.au](mailto:info@cyber.com.au) *Cybersource*

resources, but in general the direction is being set by the general Linux community.

In fact, the Fedora Project consists of a number of separate parts. The first is Fedora Core, which is the basic distribution, blessed by the steering committee. It consists entirely of free software, and will build an entire general-purpose operating system. This is augmented by Fedora Extras, which are extend but do not replace the Fedora Core, e.g. "Fedora Extras HPC" for high performance computing. There is also Fedora Alternatives, which are packages that replace parts of the Fedora Core.

Finally, and possibly one of the most important initially, is Fedora Legacy. This part of the project will provide support for old versions of core packages that have been superseded. In fact, Fedora Legacy will also be providing support for earlier version of Red Hat Linux after Red Hat's end-of-life date has been passed. While this level of support will not be the same as Red Hat may have been provided, it will be an ongoing avenue for most of the Red Hat community.

The one area that the Fedora Project does need to pickup is publicity. If their press-releases had been released at the same time as Red Hat's announcement, much of the Linux community's concern may have been lessened.

Okay, now we are passed the unpaid political announcement, and onto where this all fits into my Home Network (after all, that is the title). The Fedora Project has released "Fedora Core 1" (FC1), their first full package, and I've downloaded and installed it. Not surprisingly, it is not significantly different from Red Hat 9, and is in fact derived from the the beta version that Red Hat was developing for their next release. It even goes so far as to be code named "Severn" and includes many packages that are in other Red Hat releases (e.g. Redhat-artwork, redhat-config-bind, etc).

The main aim of FC1 was to get a product out there, but as with any update, there are a number of minor updates. The most major one is that the gcc has been upgraded to GCC 3.3. An interesting side effect of this is that Java needs to have a special plug-in compiled with gcc-3.2, such as the plugin included with Sun j2re 1.4.2, in the ns610-gcc32/ directory. If the wrong plug-in is linked in, then mozilla will not start, with no obvious error.

When upgrading from Red Hat 9 to Fedora Core 1, things proceed pretty much as with any previous upgrade with the exception that there is not a prompt to select packages (this does occur during a fresh install). In addition, while I haven't used it, there is also an option to do the install using VNC, so that it can be performed remotely. The upgrade process also removes any packages that have been dropped from FC1, for example, replacement of LPRng with CUPS.

All in all, I found that Fedora Core 1, was a simple upgrade and the look and feel was just what I'd expect from any Red Hat Linux. Most of the packages were the

same and worked as expected. Later versions of Fedora Core will be much more interesting, in that their next version will include the soon to be released Linux 2.6 kernel, and will diverge further and further.

That will do for another year, I hope you have a good Christmas and Happy New Year, and if you have any interesting items or comments, let me know.

See you all next year.

## AUUGN CD-Rs in this issue

Author: Greg Lehey <[Greg.Lehey@auug.org.au](mailto:Greg.Lehey@auug.org.au)>

As explained in the last issue of AUUGN, we had some minor problems with the CD-R we had planned for that issue: somebody had helped themselves to too many of the CD-Rs, and we didn't have enough left for the issue. That has its positive side, though: the CD-R was Knoppix 3.2, and since then version 3.3 has come out. That's one of the CD-Rs in this quarter's edition. The other one is the ever popular FreeBSD 4.9, which looks like being the last version of FreeBSD release 4.

It's only been nine months since we last included a Knoppix CD-R in AUUGN, but the response was so positive that when a new version came out, we decided to do so again.

### KNOPPIX 3.3

Knoppix (the "K" is pronounced) is a standalone Linux distribution designed to run without a hard disk. It's useful as a rescue disk or as a demonstration, since it will run on just about any computer. I tried the previous CD on my Dell Inspiron 7500 with a 1400x1050 display, which had given me problems with other systems in the past. It also had a Lucent wireless card. Knoppix recognized and configured both the display and the wireless card correctly, which I found quite impressive. On one of my development test machines it came up correctly, but of course it didn't recognize that the monitor was an ancient ICL VGA incapable of more than 640x480. Fortunately I was able to switch back to that resolution and use it anyway. At this point, though, I should point to the following disclaimer, which I have decapitalized to save pain:

*Disclaimer: This is experimental software. Use at your own risk. Knopper.net can not be held liable under any circumstances for damage to hardware or software, lost data, or other direct or indirect damage resulting from the use of this software.*

The point here is that Knoppix can't know anything about older monitors. Get things wrong and it will burn out the monitor. With a bit of finger trouble you might also find a way to overwrite the hard disk on the machine on which you're running. With a bit of care,

though, you should find it a useful tool.

You can mount the CD on another system to look at things, of course. There's some documentation in the directory KNOPPIX, in particular the file KNOPPIX/KNOPPIX-FAQ-EN.txt. There's more documentation in the directory Talks, but this most of it is in German.

To start KNOPPIX, just boot from the CD-ROM. It comes up with a functional KDE 3 desktop and doesn't use any local disk. A rather strange quirk is that it doesn't allow login at all, so you can't get a root shell. Instead use sudo without a password. See the file KNOPPIX/README-Security.txt, which is in German first and then English.

## FreeBSD 4.9

The other CD-R in this issue is the first CD of the FreeBSD 4.9 set. It contains the complete system, including sources, and it boots and installs in 32 bit Intel platforms. For other platforms, see <http://www.FreeBSD.org/> and contact Liz Carroll if you'd like a boot disk for one of them. There are complete installation instructions in the file INSTALL.TXT on the CD.

For the experienced, though, the procedure is simple. The following text is reproduced with permission from my book "The Complete FreeBSD":

- If you have another operating system on the machine, for example Microsoft, and you want to keep it,
  1. Make a backup! There's every possibility of erasing your data, and there's absolutely no reason why you should take the risk.
  2. Repartition your disk with FIPS, which is available on the CD at tools/fips.exe.
- Insert the CD-ROM in the drive before booting.
- Boot the FreeBSD system. The easiest way is to boot directly from the CD.
- Select the Custom installation: it's the only one which allows you to back up a step if you make a mistake.
- If you have repartitioned with FIPS, in the partition editor, delete only the second primary Microsoft slice. The first primary Microsoft partition contains your Microsoft data, and if there is an extended Microsoft partition, it will also contain your Microsoft data. Then create a FreeBSD slice in the space that has been freed.
- Otherwise delete whatever you may find in the partition editor and create new FreeBSD slices.
- On exiting from the partition editor, select the BootMgr MBR.

- In the disk label editor, select the FreeBSD slice. If you proceeded as above, it should be empty, but if it contains existing UNIX partitions, delete them. If you're not too worried about the exact size of the partitions, select automatically generated disk labels.
- Alternatively, if you want to specify your file systems yourself, start on the basis of a root file system with 50 MB, a swap partition with 256 MB, and allocate the rest of the space on the disk to the /usr file system. Note particularly that, if you don't create a /var file system, you'll need to create a symlink later on.
- Choose the distributions you want. Note that in this menu, you choose the distribution by pressing the space bar, not the Enter key.
- Select CD-ROM as the installation medium.
- If you intend to run the X window system, select the installation now. It's much easier than doing it after the system is up and running.
- Confirm installation. The system will be installed.



## Who Are You?

The AUUG'2004  
Annual Conference  
Melbourne, 1-3 September  
2004 Tutorials 29-31 August  
2004

## Call for Papers

As more devices, companies and people get connected to the Internet, computer security becomes increasingly important. And often security boils down to three things:

- **Identification** – working out who you are dealing with.
- **Authentication** – confirming you know who you are dealing with.
- **Authorisation** – letting the known person do what they are allowed to do and no more.

With that in mind, AUUG has chosen as the theme for the 2004 conference: "Who Are You? Identification and Authorisation Issues in Computing.", and invites proposals for papers and tutorials relating to:

- Identification, authentication and authorisation
- Applications of cyptography and cryptographic protocols
- Maintaining privacy



- Achieving anonymity on the Internet
- Internet security
- Other aspects of computer security

We also call for papers relating to topics of general interest to AUUG members:

- Standards based computing
- Open source projects
- Business cases for open source
- Open source in government
- Technical aspects of Unix, Linux or BSD
- Performance measurement and management
- Software development
- Networking, Internet and the World Wide Web.

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

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

Speakers may select one of two presentation formats:

Technical presentation: a 30-minute talk, with 10 minutes for questions.

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

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

Tutorials (held 29-31 August) provide a more thorough presentation, of either a half-day or full-day duration. They may be of either a technical or management orientation.

The AUUG'2004 conference offers an unparalleled opportunity to present your ideas and experiences to an audience with a major influence on the direction of computing in Australia.

#### SUBMISSION GUIDELINES

If you are interested in submitting a paper you should send an extended abstract (1-3 pages) and a brief biography, and clearly indicate their preferred presentation format.

If submitting a tutorial proposal you should send an outline of the tutorial and a brief biography, and clearly indicate whether the tutorial is of half-day or full-day duration.

#### SPEAKER INCENTIVES

Presenters of papers receive free registration to the conference (1-3 September), including social functions, but excluding tutorials.

Tutorial presenters may select 25% of the profit of their session OR free conference registration. Past experience suggests that a successful tutorial session generate a

reasonable return to the presenter.

Please note that in accordance with GST tax legislation, we will require the presentation of a tax invoice containing an ABN for your payment, or an appropriate exempting government form. If neither is provided then tax will have to be withheld from your payment.

#### IMPORTANT DATES

Abstracts/Proposals Due:	7 May 2004
Authors notified:	4 June 2004
Final copy due:	2 July 2004
<b>Tutorials:</b>	29 to 31 August 2004
<b>Conference:</b>	1 to 3 September 2004

Proposals should be sent to:

AUUG Inc.  
PO Box 7071  
Baulkham Hills BC NSW 2153  
Australia

Email: [auug2004prog@auug.org.au](mailto:auug2004prog@auug.org.au)  
Phone: 1800 625 655 or +61 2 8824 9511  
Fax: +61 2 8824 9522

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

<http://www.auug.org.au/events/2004/auug2004/>

## Linux.Conf.Au 2004

Author: Michael Davies <[lca2004-organisers@linuxsa.org.au](mailto:lca2004-organisers@linuxsa.org.au)>

IBM underscores Linux commitment with  
Linux.conf.au sponsorship

### REGISTRATIONS HAVE OPENED FOR LINUX.CONF.AU 2004

Organisers of linux.conf.au 2004 in Adelaide have welcomed IBM's decision to sign-on as the event's Penguin Sponsor whilst announcing the opening of the conference registrations.

Linux.Conf.Au is Australia's premier conference on the technology of the Linux system. The conference has a strong technical focus, rather than being a trade show or a marketing event.

The penguin sponsor named after the system's famous penguin mascot, is the prime sponsor of the conference.

According to Geoff Lawrence business manager, IBM Australia: "In its short history Linux has proven to be one of the most important forces for the future of the information technology industry. As businesses begin to take advantage of the Internet to become on demand businesses they need to integrate their business processes and the applications that run them. Because Linux is developed by an open community that includes some of the best programming minds in the world, many of the innovative new applications that will drive

e-business on demand will be written for Linux.

"This is IBM's third year at linux.conf.au and as we were the first major commercial vendor to embrace Linux broadly we're pleased to be involved. IBM itself participates in the growth of Linux through our Linux Technology Centre, made up of more than 250 engineers worldwide who work full-time on Linux as part of the open source community. And that's not to mention more than 7 500 IBM employees working on Linux in porting centers, research, services, development labs, and sales and marketing."

Conference organiser Michael Davies explained: "We are delighted that IBM has continued to show such a high level of support for linux.conf.au."

Linux.Conf.Au 2004 will be held at The University of Adelaide between Wednesday 14 and Saturday 17 January 2004. Prior to the conference there will be two days of mini-conferences. These mini-conferences focus on particular niche aspects of Linux, such as Linux's increasing use in education or refining Linux's Ipv6 technologies.

Registration fees for linux.conf.au 2004 will be: \$600 for professional delegates, \$275 for hobbyist delegates, and a law \$99 for students. In addition to attendance at the four days of the conference, the charges include free admission to the two days of mini-conferences.

Linux.Conf.Au 2004 organisers Michael Davies said: "We want linux.conf.au to be affordable to all, especially students and hobbyists. We've made special efforts to keep the cost of the registrations as low as possible whilst still offering a great selection of speakers at the conference.

"The conference has a good selection of international speakers and the response to the Call for Papers are showing that the technical programme will be very strong."

The most recent linux.conf.au was held in Perth and had over 400 delegates. International speakers at that conference included Linux creator Linus Torvalds and renowned Linux kernel programmers Alan Cox and H. Peter Anvin.


More information and on-line conference registration can be found on the web at:

- <http://lca2004.linux.org.au/>

# Collaborative Virtual Workspaces and their advantages

Author: Carlos Andrés Pérez <[caperez@usaca.edu.co](mailto:caperez@usaca.edu.co)>

Translated into English by: Juan Manuel Triana



## ABSTRACT

This document is the product of the experiences obtained in the development and execution of the first course in college teaching from the Santiago de Cali University and other courses given via collaborative virtual workspaces, videoconference and chat. The foundation of these courses began with the following problem: Which are the conditions, criteria and methods required for the development of learning scenarios based in technological mediation, that favor new learning ways and new relations between teacher and student? The article is the result of this first learning experience with collaborative virtual workspaces Colombia.

## INTRODUCTION

Right now the virtual education gives us, among many things, immediacy, telepresence and interaction and this is different from the conventional model for higher education.

Several learning theories could be applied into this kind of virtual environments, for example Piaget, Vigotsky [2] and Dewey. The collaborative learning environments prepare the student to: [1,4]:

- Actively participate in the collective construction.
- Integrate into a virtual team or group.
- Give help to others and ask for it when needed.
- To provide its own knowledge and strength to others.
- To understand other people needs.
- To discover solutions that are good for all.
- Make significant contacts with communities from different cultures.
- To compare beliefs and activities with the other ones.
- Establish goals, tasks, resources, roles, etc.
- Listen with a critical attitude and respect the others.
- To expose own ideas and thesis using argumentation.
- Accept the reasonable critic from other people.
- Provide credit to other people.
- Negotiate language and methods.

- To develop interpersonal abilities.
- Get familiar with democratic processes.

The Internet's increased growth has resulted in creation of cooperative and collaborative environments that cross borders. Now the students can go outside from their everyday world to get involved in adventures with partners from every part of the world. This kind of work, developed by the net, is made of projects in the form of didactic experience that must be done in groups, where the participants are not a mere pen pals, but a team searching for a goal. In the planning of these activities is very important to have effective communication, so communication and interactivity are necessary and fundamental [3].

Every day the learning multimedia applications are more rich and integrated, also the powerful platform of the Communication and Data Network in diverse formats has been applied for the Grupo de Investigación en Educación Virtual de la Universidad Santiago de Cali (GIEV: <http://cvw.usaca.edu.co>) Virtual Education Research Group in the creation of pedagogic environments. It is known for its versatility, handling and delivering of information available in electronic format. These environments favor the learning process and the research of the students. One of the goals of this research group is to guide the University directors and teachers for a better management of the technical resources, with a broad and clear vision of their possibilities and role in the educative environment.

All this thinking research is based in the belief that the technology alone isn't capable to bring the quality in an educative system, but its planned use and its critical adoption based characteristics and the lessons derived from the experience and the research in different contexts.

The Virtual Education can't be understood only as a method for distance education or associated with programs in that format, or only for higher education. The GIEV is checking the possibilities to develop new technology-based mediation scenarios to grow the alternatives for the present curriculum with new virtual environment that provides for the extra value to the learning models at different levels.

With these lines, we began a series of actions to implement the virtual education model, at first, inside our institution, because of the viability to coordinate efforts and human, technical and logistical resources for a task that could generate important changes in the everyday educative activities and that could be projected from the GIEV to the main academic community. The Virtual Education Project is not only providing the possibility to generate new teaching forms, but to improve a process that is been changed in the last ten years, mainly because of the adaptation of new technologies with a growing influence in the society.

## MATERIALS AND METHOD

The problem base in order to formulate this work was

conceptualized this way:

Which are the conditions, criteria and methods required for the development of learning scenarios based in technological mediation, that favor new learning ways and new relations between teacher and student?

Inside our first approaches with the directives we found a broad range of interrogations that was distilled during two years of meetings and tests:

- What's the meaning of Teaching in Technology?
- Which are the bases to incorporate media in the curriculum from a pedagogic view?
- Which are the existent relation(s) between the teachers and students in an university, with the technological resources for communication and information?
- Which scientific, pedagogic and computer aspects are required to get the appropriated incorporation of technology in educative processes?
- Which scientific, pedagogic and computer aspects are required to implement a Virtual Education Model in the university?
- How these technological devices influence the learning and teaching dynamics?
- When we speak about Virtual Education are we thinking of it as a Distance Education?

The decision of research in the field of technology was applied to the curriculum used inside the Santiago de Cali University. The idea is now that, over time, the teachers, scholars and alumni from different levels could develop and apply analytic abilities to the use of new instruments and the communication channels.

The model chosen for the course was that of Collaborative Learning. The technological platform was based on Open Source Tools. The Operating System used has been Linux, that has responded with great stability, speed and reliability as server.

We used as main software Claroline (<http://www.claroline.net/>), UNESCO approved and open source, made in PHP, which permit an easy adaptation, complement it or use it as the model for further development.

The databases has been handled using MySQL, for a daily average of 80 users at the same time.

The hardware used has been very cheap, because Linux can be implemented in PCs, this has let that the institution uses its own technical resources in this earlier steps (planning, development, test), without spending money on expensive servers.

## RESULTS

Using the collaborative virtual workspaces, software tools and the collaborative learning model we achieved the following:

- Eliminate one of main problems in Distance

Education: the student isolation. Based on the professional experience in this field, we can affirm that the distance model has few moments of true contact and interaction between the student groups and the teacher or tutor. The new technologies generate important changes in concepts like distance, time and communication forms.

- Students and Tutor-Teachers build together the knowledge and they develop tolerance and cooperation, necessary in a world where the knowledge is every day more socialized.
- The computer world has and enormous impact on the younger population and they feel that technologies are part of their environment and, specially, are part of their future work environment. The virtual education model, through networks, could transform a computer screen in a new open window to collaborative learning.
- It could help to solve the main problems (production and distribution of print and audiovisual material) that have been one of the main factors that caused in the past distance education programs to fail. The distributed information is not limited to the institution or teacher proposals because the students can access in the network lot of data and multiple information sources.
- The flexibility of the virtual education model, because it can be adapted to the learning needs of the diverse target groups. For that reason the best model to virtual education and distance education is the collaborative work in virtual learning media, because it stimulates the individual and group participation.
- The viability to conduct an historical record of all the learning development process and the interaction between students and the teachers. This facilitates the design of new pedagogic tools based on the previous course results. This has been absent in the traditional distance education.
- Promote, using chat, the organization of ideas in a written way.

These and another reasons assure us that the college processes in virtual and distance education can be extended to many national and international places.

## BIBLIOGRAPHICAL REFERENCES

- [1] Manuel Antonio Unigarro Gutiérrez, Educación Virtual: Encuentro Formativo en el Ciberespacio. Editorial UNAB, Bucaramanga, Colombia, 2001.
- [2] Luz Adriana Osorio Gómez, capítulo "Aprendizaje en Ambientes Virtuales y Colaborativos" del libro "Los Computadores en la Nueva Visión Educativa", Escuela Colombiana de ingeniería, 2000.
- [3] Propuesta de Integración de las Tecnologías de Información y Comunicaciones a los Centros Escolares de Fe y Alegría. <http://www.fevalegría.org/>
- [4] José Guadalupe Escamilla, "Selección y Uso de Tecnología Educativa", segunda edición, Trillas, ITESM, 1999.


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

<http://www.linuxfocus.org/English/September2003/article312.shtml>

# A Short Guide to the Installation of CVW Under Linux

Author: Juan Manuel Triana <[jutriana@uniweb.net.co](mailto:jutriana@uniweb.net.co)>  
and Carlos Andrés Pérez <[caperez@usaca.edu.co](mailto:caperez@usaca.edu.co)>

Translated into English by: Juan Manuel Triana



## INTRODUCTION

CVW (Collaborative Virtual Workspaces) is one of the finest open source tools available for the development of virtual education. This article describes the process of installation and running in Suse Linux Professional 8.0.

Further development on the original Cvw is not taking place at this moment but the software is currently being translated to Spanish and is of course used in the courses at Universidad Santiago de Cali in Colombia. As of the writing of this article there are about 600 users of Cvw at the university.

## WHAT IS CVW?

CVW (<http://cvw.sourceforge.net/>) is a collaboration software environment written mainly in Java, that provides a "virtual building" where teams can communicate, collaborate, and share information, regardless of their geographic location. With CVW the interactive possibilities of the net can be used in virtual spaces where the people gathers to learn, interact, discuss, share information, using a series of friendly tools. CVW has been developed by the Mitre Corporation (<http://cvw.sourceforge.net/>) and is completely Open Source.

## WHY WE CHOOSE CVW?

The GIEV (<http://www.gievirtual.com/>) group at Universidad Santiago de Cali in Colombia, has been working for several years in researching new platforms and virtual environments that allow an easy integration with the existent education methods. After several testing of commercial software we began to look for in Open Source development and discovered CVW. It has all the main elements that we were looking for, so we decided to try.

CVW is not easy to install, as we soon discovered. In the next lines I'll describe the process to put running this wonderful platform on Suse 8.0.

We used an IBM Netvista machine to use as server for this software. The machine (<http://cvw.usaca.edu.co/>) has a Pentium 4 at 2.0Ghz, 256MB Ram and 40G Harddrive, connected to the University network. The Suse installation went clean, only the video give me some trouble (an integrated Intel 845), but with a kernel compilation the video was soon running well (1024x768 at 24 bits).

## STEP BY STEP

The first step is to download the software from the project website (<http://sourceforge.net/projects/cvw/>):

We need to get the following applications:

CVW server  
CVW document server

There is also the CVW Federated server available which applies only to a cluster of CVW servers.

First install the CVW server. We untarred the file in /opt/CVWserver and created an user to run it. The script to run the server is

```
./cvw.boot start
```

And to end the execution:

```
./cvw.boot stop
```

The software recommended to create an user, we created the user cvw and gave it the rights on the directory. When you look at the script file (cvw.boot) you must enter some data like the installation directory and, the most important, the TCP port where CVW will check for client connections. By default the assigned port is 8888. You only have to add it in the /etc/services script. The only thing that you must change when installing under Linux is the line that fixes the maximum number of clients, by default is unlimited, but in Linux this will cause the service to crash.

When you connect with a browser you will be greeted with the following message:

```
<!--YourMOO-->
<!--
Welcome to the Collaborative Virtual Workspace.
This is version 4.0.2 of the CVW core.
If you are not using a CVW client ...
type: connect username password ...
to connect as an existing user
type: quit
to disconnect
-->
I don't understand that. Valid commands at this
point are
Help, Connect, or Quit
```

This means that our server is ready to serve any client. The second step (and the most painful) is to install the document server.

Untar it in /opt/cvwdocserver and look at the README

file. Now it is time to make your first decision.

The CVW docserver uses MySQL as database and Tomcat (<http://jakarta.apache.org/>) (an Apache Java utility) to handle all documents stored in the virtual workspace. The tar.gz file includes a jakarta/tomcat directory and a complete mysql installation (3.23.27-beta)

We don't like beta versions and Suse had a newer jakarta and mysql versions that were running fairly well, so we decided to adjust it to the docserv requirements. With mysql there were no problems, we only move the database file to my existent installation (/var/lib/mysql).

The real problem came with the Tomcat configuration. Trying to adapt the existent jakarta to the configuration of CVW's tomcat was a headache, because the commands were located in different parts, so at least we decided to run the included Jakarta/Tomcat.

But you must change several thing in order to get a running application: we left the jakarta/tomcat directory under my docserv tree (/opt/cvwdocserv). In your apache configuration file (/etc/httpd/httpd.conf) you must add a line pointing it:

*(Editor's Note: the following is all on one line in the apache configuration file)*

```
Include
/opt/cvwdocserver/jakarta/tomcat/conf/tomcat-
apache.conf
```

Now to the tomcat-apache.conf file. You have to give the exact address of the jserv module in apache:

```
LoadModule jserv_module libexec/mod_jserv.so
```

Also you must indicate the correct address of your tomcat installation and the port to listen (my choice was 8007, after several trials):

```
ApJServDefaultPort 8007
AddType test/jsp .jsp
AddHandler jserv-servlet .jsp
Alias /examples
/opt/cvwdocserver/jakarta/tomcat/webapps/examples
<Directory
"/opt/cvwdocserver/jakarta/tomcat/webapps/examples"
>
Options Indexes FollowSymLinks
</Directory>
ApJServMount /examples/servlet /examples
<Location /examples/WEB-INF/ >
AllowOverride None
deny from all
</Location>
Alias /test
/opt/cvwdocserver/jakarta/tomcat/webapps/test
<Directory
"/opt/cvwdocserver/jakarta/tomcat/webapps/test">
Options Indexes FollowSymLinks
</Directory>
ApJServMount /examples/servlet /examples
<Location /examples/WEB-INF/ >
AllowOverride None
deny from all
</Location>
Alias /test
/opt/cvwdocserver/jakarta/tomcat/webapps/test
```

```

<Directory
"/opt/cvwdocserver/jakarta/tomcat/webapps/test">
Options Indexes FollowSymLinks
</Directory>
ApJServMount /test/servlet /test
<Location /test/WEB-INF/ >
AllowOverride None
deny from all
</Location>
ApJServMount /servlet /ROOT

```

Once this is done, you must change the docserver script, cvwds.boot that use the same arguments as cvw.boot (start | stop). We only changed the local directories for the docserver and the database and commented the line that shuts down the mysql server (because we have our own databases running).

Also you must check that the .jar files inside the webapps/app directory must be there and not in subdirectories.

If you are lucky you are almost done.

The final step is to untar the client file. We did it under /opt/CVWclient/. The script with all the needed data is called server.cvw, in my case the lines that we modified were the following:

```

# Property file for CVW
#
cvw.server.host=cvw
cvw.server.port=8888
cvw.server.name=cvw
cvw.docserver.host=cvw
cvw.docserver.port=8080
cvw.docserver.url=/app/docervlet
cvw.version=4.0
cvw.debug=true

```

```

#####
# User Images
#####
#this url must end with a /
cvw.userimages.url=http://localhost/cvw/user-
images/

```

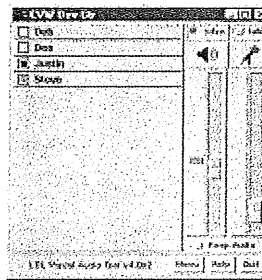
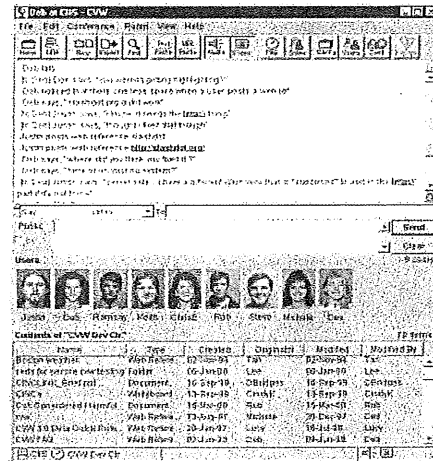
The cvw.docserver.url is key... we don't know why the original points to another direction and you discover after several trials that this is the MOST IMPORTANT line to get acces to your docserver.

At this moment you are ready to enter the first time to your CVW server. For your first login you will enter as "Admin" with no password.

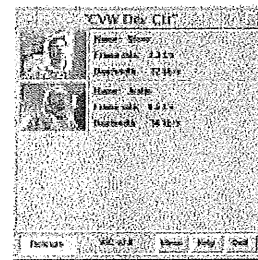
Now you can change your password under "File" and use the Admin Tool under "View".

As the menu shows, there's a lot of tools available and that make this application a reference in virtual education.

CVW Client Interface




Audio Conferencing

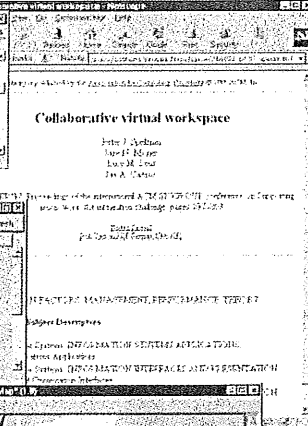


Video Conferencing



Private Data




Shared Web Browsing



Online Users

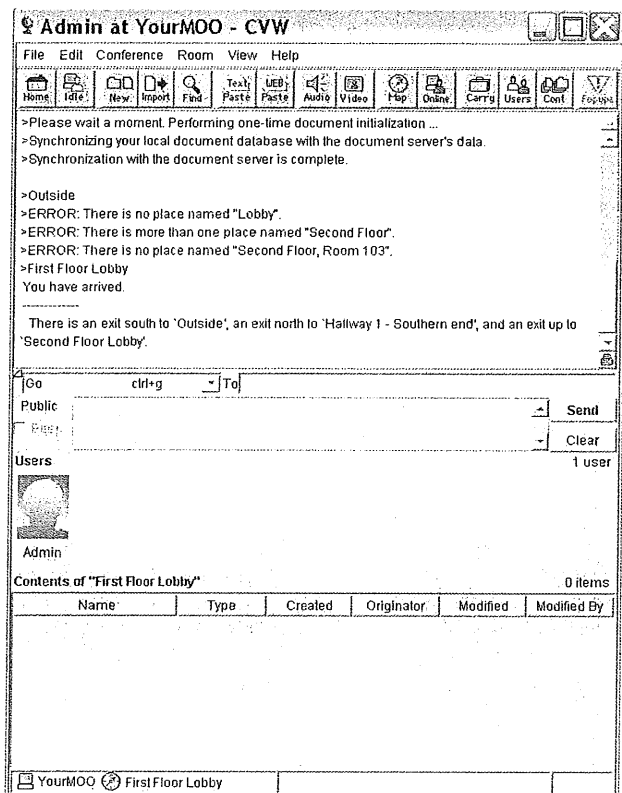


Virtual Building Floor Plan



Shared Whiteboard

This configuration is working now in our server. You can try this utility. If you wish want information: they



have a good manual in pdf.

Enjoy.

...

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

<http://www.linuxfocus.org/English/September2003/article310.shtml>

## Building 'Nefarious' 3/5

Author: [pasnak@warpedsystems.sk.ca](mailto:pasnak@warpedsystems.sk.ca)

### INTRODUCTION

As I stated in the previous article, this section will cover the customization and benchmarking of Mandrake Linux 9.2 AMD 64. I'll be covering the NVidia AMD 64 driver installation, configuring DVI, testing the DVD burner, some system tweaks and some benchmarks (hdparm, oggenc and kernel recompile). I've had to recompile some apps from Mandrake and PLF, so I'll provide whatever insight I can as to how I got them to work. The biggest thing I have noticed with building Nefarious is that as a rule of thumb the only time something didn't work exactly the same as it did on x86, or something didn't work at all, it was closed source.

### NVIDIA DRIVER/DVI

I know some of you have strong feelings towards closed drivers, but nobody is coming over to your house and forcing you to install them (yet), so learn to relax a

little.

The first thing to do is to download and install the kernel-source RPM that matches the running kernel. In this case it is kernel-source-2.4.22-24mdk:

```
$ wget ftp://ftp.sunet.se/pub3/os/Linux/distributio
ns/mandrake-devel/cooker/amd64/Mandrake/RPMS/kernel-
source-2.4.22-24mdk.amd64.rpm
```

```
$ su
# rpm -ivh kernel-source-2.4.22-24mdk.amd64.rpm
```

Next we grab the latest NVidia drivers for AMD 64 from <http://www.nvidia.com/content/drivers/driver.asp>. In this case it is the 1.0-4499 driver. They don't use the new install method that is available for the x86 platform, instead they use the old GLX and kernel tar.gz files. Not a big problem, but worthy of a note. So to get the driver ready, we do the following:

```
$ wget http://download.nvidia.com/XFree86/Linux-
x86-64/1.0-4499/NVIDIA_GLX-1.0-4499.tar.gz
$ wget http://download.nvidia.com/XFree86/Linux-
x86-64/1.0-4499/NVIDIA_kernel-1.0-4499.tar.gz
$ tar -xvzf NVIDIA_GLX-1.0-4499.tar.gz
$ tar -xvzf NVIDIA_kernel-1.0-4499.tar.gz
$ su
# cd NVIDIA_GLX-1.0-4499
# make
# cd ../NVIDIA_kernel-1.0-4499
# make
```

Both compiled okay, and the NVidia kernel driver was loaded upon completion of the compile. I logged out of KDE, switched to a console as root, and disabled the X server:


```
<CTRL-ALT-F1>
$ su
# init 3
```

Editing '/etc/X11/XF86Config-4', I find that the only thing I have to change is the name of the driver from "nv" to "nvidia". Everything else is already there (load "glx", etc), and the XFree86 logs show that these settings are just skipped over if no 3D capable driver is found. If you run 'XFdrake' after making your changes, it will honour the 'nvidia' selecting, but also add 'alias /dev/nvidia\* nvidia' to '/etc/modules.conf', using that the driver is loaded during boot (or you can add it manually).

I switched back to runlevel 5 (init 5) to restart X, and all that I saw was a black screen with a blue 'smudge' at the bottom. The system was unresponsive, so I logged in remotely via openssh, and noticed that the X server was taking up 99% of the CPU. I rebooted the system into runlevel 3, disabled the nvidia driver and started to hunt down the problem. After googling my brains out, I finally discovered that it might be an issue with AGPGART, so I set XF86Config-4 to use NvAGP by default (details below), and the NVidia splash screen popped up on the next reboot.

'glxgears' output, although impressive when compared

to my old system, was not earth-shattering. All the OpenGL screensavers worked also.



```
$ glxgears
7289 frames in 5.0 seconds = 1457.800 FPS
8537 frames in 5.0 seconds = 1707.400 FPS
8523 frames in 5.0 seconds = 1704.600 FPS
8405 frames in 5.0 seconds = 1681.000 FPS
8244 frames in 5.0 seconds = 1648.800 FPS
```

I'll see if I can tweak this to provide better output later on. It was working though, so getting DVI to work was a fairly simple process. Following the direction in the NVidia README, I added the 'ConnectedMonitors' option, then shut the system down and switched the connector over to DVI. Initially, I could only get digital output from the console, so I took a look at '/var/log/XFree86.0.log' to see if I could determine what the problem might be. The first thing I noticed was that it was referencing the CRT as the primary device, regardless of the order I put the options in for 'ConnectedMonitor'. Removing 'CRT' from the option let it continue, but lead me to my second issue - resolution size. The 'XFree86.0.log' output was very detailed and easy to understand - I had to switch to 1280x1024 if I wanted to get everything to display properly. Everything is slightly 'bigger' than I am used to, but it is very crisp, so we'll leave it as it is for now (I was able to increase the framerate by about 300fps in glxgears by setting 'RenderAccel'). Below is the relevant section from XFConfig-4 (/etc/X11/XF86Config-4):

```
Section "Device"
Identifier "device1"
VendorName "NVIDIA"
BoardName "NVIDIA GeForce FX (generic)"
Driver "nvidia"
Option "DPMS"
Option "ConnectedMonitor" "DFP"
Option "RenderAccel" "1"
Option "NvAgp" "1"
EndSection
```

There must be something mystical about the 'NvAgp' option, as it appears to become invisible when typed into any support/discussion forum. No matter what forum I went to, it was one of the first things mentioned as a possible solution to peoples problems, but in 99 out 100 of them, people seemed to ignore this option, and opted to try rebooting a couple dozen times instead. Either that or copy the entire contents of XF86Config-4 to the forum and complain that SUSE/Mandrake/Gentoo/Slackware/whatever sucks. I don't claim to understand, I'm just telling you what I read.

## SYSTEM TWEAKS

Before I moved on to the DVD Player, I figured I'd tweak the system a bit. I started by installing lineakd to take advantage of all the keys on my keyboard that after next week will probably never get touched again -- except for the little scroll wheel on the keyboard, that kicks ass -- and yes, since my keyboard wasn't listed, I created a template and submitted it to the author of lineakd.

Next I grabbed some icon sets from KDE Look and some desktop backgrounds from Digital Blasphemy. I settled on Korilla and Tropical Moon of Thetis. I also grabbed some fonts from the web, and used 'drakfont' to add them to the system.

The next thing I wanted to get going was video playing. MPlayer was already available for AMD64, but I'm a sucker for movie trailers, so unfortunately, I would need the ability to play WMV files. The win32-codecs where available from PLF, but not for AMD64. I grabbed the source package (win32-codecs-1.4-2plf.src.rpm), installed it and made the following edits to win32-codecs.spec (in diff format):

```
16c16
<
---
> ExclusiveArch: %ix86
```

After that was done, I did a 'rpm -bb win32-codecs.spec' then installed the resulting '.amd64.rpm'. I was expecting problems running .wmv, but everything I downloaded worked, and I was also able to play .asx streams. I didn't have any luck compiling kmpayer (there is a few hours of my life I'll never get back...).

By this time the onboard soundcard was really starting to get on my nerves. After setting everything (mplayer, xine, etc) to use artsd for output, it worked a little better, but it was stile sluggish when changing songs -- about 1/2 second behind. So I pulled out an SB Live! Value I had lying around and plugged that in. I ran 'alsaconf' as root to configure the modules for the card -- which promptly overwrote my '/etc/modules.conf' (all hail backups). The card sounds much better, and is far more responsive.

```
/etc/modules.conf

alias eth0 sk98lin
probeall usb-interface usb-uhci ehci-hcd
alias ieee1394-controller ohci1394
probeall scsi_hostadapter sata_promise sata_via id
e-scsi
alias /dev/nvidia* nvidia
# --- BEGIN: Generated by ALSACONF, do not edit.
# --- ALSACONF verion 0.9.0 ---
alias char-major-116 snd
alias snd-card-0 snd-emu10k1
alias char-major-14 soundcore
alias sound-slot-0 snd-card-0
alias sound-service-0-0 snd-mixer-oss
alias sound-service-0-1 snd-seq-oss
alias sound-service-0-3 snd-pcm-oss
```



```
alias sound-service-0-8 snd-seq-oss
alias sound-service-0-12 snd-pcm-oss
options snd major=116 cards_limit=1 device_mode=06
66
options snd-emu10k1 index=0 dxs_support=2 snd_exti
n="0x0003" snd_extout="0x1f0f"
# --- END: Generated by ALSACONF, do not edit. ---
```

I continued to tweak the soundcard, using 'alsamixer' to enable digital output, center speaker, etc.

The last thing I wanted to get going was kopete, which is a flexible and extendable multiple protocol instant messaging system designed as a plugin-based system -- Came up with that off the top of my head. OK, I didn't, it's what their web page describes it as. In short, kopete is ICQ, AIM, MSN, Jabber, etc all wrapped up into one interface, but at this point it wasn't available as a '.amd64.rpm'. I grabbed the source (kopete-0.7.3-0.92mdk.src.rpm) and tried to rebuild it with 'rpm --rebuild kopete\*.src.rpm'. It didn't work the first go, so I had to make some changes to the spec file (in diff format):

```
kopete.spec
108a109,110
> --with-qt-libraries=/usr/lib/qt3/lib64 \
> --with-qt-includes=/usr/lib/qt3/include \
```

After this, it compiled and installed OK, so I copied over my config from my old box and fired it up.

Since I would be using Nefarious as my day-to-day work machine, I copied over necessary parts of my home directory (Mail, Documents, etc) and dropped them into my new home. Everything from bookmarks to Mail to shortcut settings are now exactly the same as my previous box. Slick.

## DVD BURNER/PLAYER

I initially thought I would experience a lot of problems with the DVD devices, but as luck would have it, it is probably the smoothest running device in Nefarious. To start out, I set both the DVD Burner and DVD Player to 'ide-scsi' -- I don't know why this just isn't the default in distributions -- in '/etc/lilo.conf', run '/sbin/lilo' and reboot:

```
image=/boot/vmlinuz
label="linux"
root=/dev/sda1
initrd=/boot/initrd.img
append="devfs=mount hda=ide-scsi hdb=ide-scsi hda1
un=0 acpi=off splash=silent"
vga=788
read-only
```

Once the system was back up, I installed k3b and k3b-plugins. Although k3b was available as a '.amd64.rpm', k3b-plugins wasn't. I grabbed the source RPM (k3b-0.10.1-2mdk.src.rpm) from MandrakeClub, and it rebuilt without error. I was able to burn CDs and DVDs, burn ISOs and create audio CDs on the fly from ogg files. If I wasn't straight forward enough in the first article - if you haven't tried k3b, get it. Use it. Love it.

Next up was DVD playing. Thankfully, I'm Canadian, so I'm not going to go to jail (yet) for wanting to watch my favourite DVDs while I code. A quick trip to The Penguin Liberation Front had me rebuilding libdvdcss-1.2.8-1plf.src.rpm for Nefarious. Now the search was on for a decent DVD player. Although MPlayer is perfect for playing movie files, I figured I'd find something dedicated to DVD playing. I decided to go with Ogle, and was quite pleased. The source file would not rebuild (2 and 1/2 pots of coffee later, I decided that the Ogle source RPM was plotting against me, and decided to get rid of it before it became fully self aware), so I built it from 'plain' source. So in approximately 27secs and a 'ogle /dev/scd1' I was quite happily watching the first season of Angel. Everything worked, from loading the libdvdcss to menu navigation to subtitles.

## LIES, DAMN LIES, AND BENCHMARKS

I'm not a serious gamer, overclocker or point-release slave, so benchmarks usually mean very little to me, but it gives me an excuse to use the chart function in OpenOffice, so what the hell. I used a dual AMD Athlon (an Altus 140 from Penguin Computing) my old AMD Athlon 1G as comparisons. All of the boxes are using Mandrake 9.2, with no major optimizations.

AMD Athlon 64 3200+ (nefarious)  
Linux Version 2.4.22-24mdk  
One 2GHz UNKNOWN Processor  
1GB RAM  
3997.69 Bogomips  
120Gig 7200RPM Maxtor SATA

Dual AMD Athlon 2800+ (brain)  
Two 2.13GHz AMD Athlon Processors  
Linux Version 2.4.22-24mdkenterprise  
1GB RAM  
8519.68 Bogomips  
36Gig 15000RPM Seagate SCSI

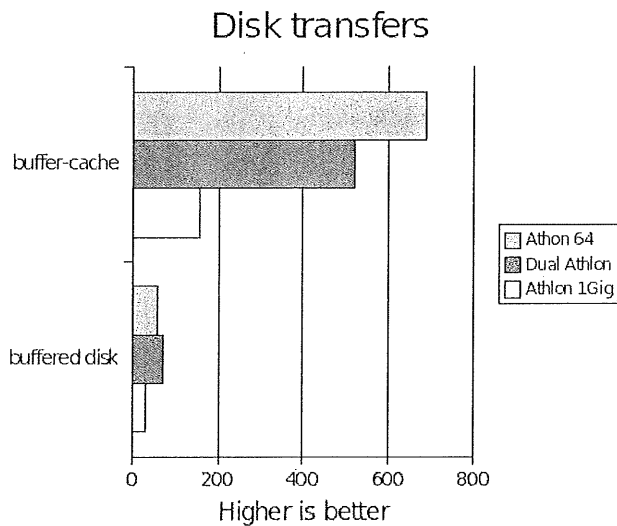
AMD Athlon 1Gig (neo)  
Linux Version 2.4.22-24mdk  
One 1.01GHz AMD Athlon Processor  
640M RAM  
2011.95 Bogomips  
40Gig 7200RPM Maxtor ATA-133

The first bench I did was encoding a wav file to ogg, using oggenc from vorbis-tools-1.0-7mdk. I selected The Mercury March, the official march of the Canadian Communications and Electronics Branch. The original file size was 37.89 Meg, with a playing time of 3m 34.0s. It was encoded with no flags, using 'oggenc Mercury\_March.wav -o mercury\_march.ogg'.

	<i>Athlon 64</i>	<i>Dual Athlon</i>	<i>Athlon 1Ghz</i>
Seconds	11.8	15.9	81.6
Rate	18.18	13.52	2.63

The next benchmark used hdparm to test the buffer-cache reads and the buffered disk reads. If you currently use a system with SCSI, you'll be happy with the overall performance of the SATA drives. The test was done with 'hdparm -tT /dev/xxx'.

	<i>Athlon 64</i>	<i>Dual Athlon</i>	<i>Athlon 1Ghz</i>
buffer-cache	690	520	156
Buffered-disk	56.48	71.33	30.62



The last bench I did was a kernel compile. This was not a vanilla kernel compile, but a Mandrake kernel SRPM rebuild. I edit the kernel.spec file of kernel-2.4.22.24mdk-1-lmdk.src.rpm to have it only produce kernel-2.4.24, kernel-source-2.4.24 and kernel-doc-2.4.24. The test was done with 'time rpm -bb kernel.spec'

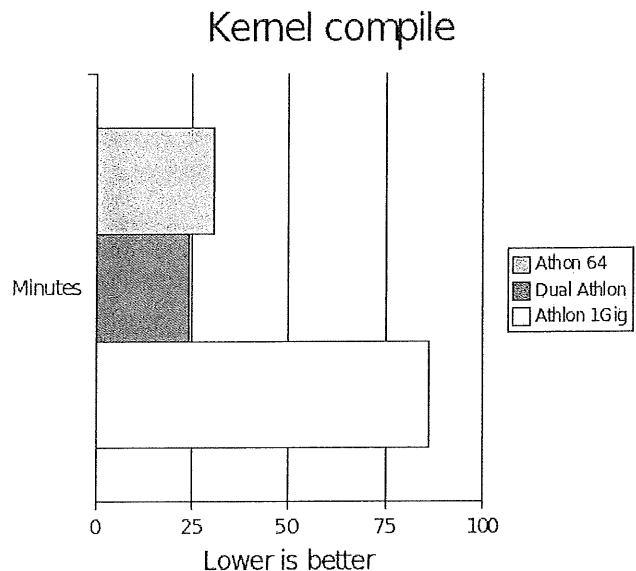
	<i>Athlon 64</i>	<i>Dual Athlon</i>	<i>Athlon 1Ghz</i>
Minutes	30.5	24.25	86

I wanted to do some 3D benchmarks, but I couldn't get anything to run reliably on all three platforms. So to sum up, the AMD 64 performs well against a Dual Athlon, and kicks the living crap out of a 3 year old AMD Athlon 1Gig. But how much stake do you put in benchmarks? If you're me, not much. You can see the full specs and text output of the tests in bench.txt.

## CONCLUSION

Overall, I'm quite satisfied with Nefarious in its current configuration with Mandrake 9.2AMD64. What you can take away from these articles, besides knowing what my cat looks like, is that the AMD 64 build will run the majority of the Linux software that is out there without problems, and run it well.

There were a few things I couldn't get running at the moment - games being one of them. Everything out




there would appear to be using the Loki installer, and it errors out with 'This installation doesn't support glibc-2.1 on Linux / x86\_64'. If anyone knows of a way around this, drop me an email. The other thing I couldn't get to run was the Flash plugin. I didn't expect it to, but I tried, and failed.

Some say the true value of the AMD 64 will not be seen until software is optimized to take advantage of the features of the AMD 64 architecture, but I don't see it that way. The true value in the AMD 64 is, as these articles have shown, the benefit of not having to throw everything you know away, and just run the software you were running before. If you are in the market for a new computer, getting a AMD 64 based system is the logical step.

In closing, I'd like to give credit where credit is due:  
 Screenshots: KSnapshot  
 Image Editing: Gimp  
 HTML Editing: Quanta+ and vim  
 Graphs: OpenOffice.org

## Ogg Encode



Hardware: EGlobalOnline and NCLIX (Dual Athlon from Penguin Computing)

Software: Mandrake Linux and the thousands of other Open Source developers out there.

I'd also like to thank my wife for putting up with me. Not just during the writing of this article, but always.

Questions/Comments can be left in this (<http://www.warpedsystems.sk.ca/modules.php?op=modload&name=News&file=article&sid=1221&mode=thread&order=0&thold=0>) thread, or drop me an e-mail.

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

<http://www.warpedsystems.sk.ca/modules.php?op=modload&name=Sections&file=index&req=viewarticle&articleid=36&page=1>

## 2003 And Beyond, Pt 2.

Author: Andrew Gryusr [aax@aaxnet.com](mailto:aax@aaxnet.com)

Editor's note: This is the second part of a series of articles which will be printed in AUUGN over the coming few issues.

### PALLADIUM

Palladium, a chip based "security" initiative, is another major part of Microsoft strategy. Since Microsoft has seized control of PC design from Intel and the PC manufacturers, they are in a position to dictate how PC hardware will integrate with Windows.

Microsoft promotes Palladium as a boon to user privacy, security, and a stopper of worms and viruses. They claim it has nothing at all to do with unpopular **DRM** (Digital Rights Management), which prevents playing, displaying or copying copyrighted content (music, videos, documents, etc.) on a PC.

That is a lie. **Truth** is found in Microsoft "help wanted" ads: "Our technology allows content providers, enterprises and consumers to control what others can do with their digital information, such as documents, music, video, ebooks, and software. Become a key leader, providing vision and industry leadership in developing DRM, Palladium and Software Licensing products and Trust Infrastructure Services" (W7).

You can be sure the primary "digital rights" Palladium protects will be Microsoft's. You will find it impossible to run Microsoft software on any computer other than the one to which it was originally registered - or anything else Microsoft doesn't want you to run.

Palladium requires computers to have a special chip, which both Intel and AMD have already agreed to incorporate. "Protected content" will not run on a computer lacking the chip, or with the Palladium features turned off. You will have to **buy all new computers** to run Palladium enabled software.

Of utmost interest is Microsoft's statement that Palladium will allow creation of **content that has an expiration date**, and which cannot be used or viewed after that date. We will look at this again in the section on Licensing and Ownership.

Palladium has been highly controversial from the very first announcement (W3), because many, including many security specialists, feel it allows Microsoft excessive control over what software you will be able to run on your PC, and the company has already stated they will not allow it to be ported to non-Windows platforms. Academics fear it will hinder the flow of ideas and destroy the doctrine of "fair use" (W24)

In fact, Palladium has become so controversial, Microsoft has done exactly what they always do when the image of one of their products becomes tarnished, they changed the name. (W32). It's now "**Next-generation Secure Computing Base**". The name has obviously been chosen to be unsuitable for slogans and titles of articles such as this one. Unfortunately for Microsoft, it's just too cumbersome, so everyone's going to stick with "Palladium".

Palladium is not, of course, technology Microsoft originated. As usual, they have commandeered technology developed by others (W4) and reinterpreted it in a way that gives greater advantage to Microsoft.

The one thing that can slow Palladium is massive consumer resistance to the way it prevents "fair use" (and well as outright theft) of copyrighted entertainment content (videos, music, games). There will be no government resistance to Palladium. Your elected representatives represent money, and their votes have already been counted.

Microsoft is already hard at work candy coating this bitter pill to get you to believe it's "good for you" (W26). Will consumers resist? Some evidence says they will. Warner Music Group, Sony Music and Universal Music

Group have already bent to consumer demand by putting a lot more titles up as singles, for download, free of "protection" schemes, and at lower prices than they original planned.

On the other hand, consumers have shown they will accept any level of abuse to avoid thinking about alternatives to Windows, so use of Palladium to protect Microsoft software (and repress competing software) is a given. Control of all other software, media and documents will follow, along with the end of free Web content in general.

## LICENSING AND OWNERSHIP

You pay plenty for Microsoft software, but you do not own it - you have a non-transferable license to use the software as Microsoft sees fit. Microsoft owns it and they are not at all shy about exercising their "property rights". Further, the license terms **can be changed retroactively** (and are) any time Microsoft pleases. It says so right in the license terms.

When Kmart filed to sell their Internet unit, bluelight.com, Microsoft told the bankruptcy court the **sale could not proceed** because bluelight.com held Microsoft property (software licenses), which could not be transferred to a new owner **without Microsoft permission**, (L14),

The situation is far worse for others. Bluelight.com had only a dozen licenses and weren't on License 6. Lets say your company is in difficult economic times and wants to sell a division with 6000 Windows workstations and some servers. You've gone with License 6 "Software Assurance" because that's the only support contract Microsoft offers. First, you have to **buy out your three year commitment** to make your licenses "perpetual". This means paying up to two years more of an exceedingly expensive support contract (25% to 30% **per year** of the cost of all your Microsoft software).

It gets worse from there - all this money benefits **neither your company nor the purchaser**, only Microsoft. License 6 **Software Assurance is not transferable** and your contract is terminated upon payment. The purchaser now needs to negotiate his own License 6 agreement starting from nowhere. To read more about this extortion, see Ed Foster's InfoWorld article (L16).

Microsoft has long chafed under the traditional PC software license, because once they sell a license, they can only keep revenue flowing by releasing "upgrades" and convincing people the upgrades are worth purchasing. Producing these upgrades is a lot of work, and they are finding it increasingly difficult to convince anyone the upgrades are worthwhile.

The solution to this problem is clearly to force users to a subscription basis, because under subscription they can continue to enjoy steady income from products while not needing to upgrade them at all. In the past, Microsoft's lust for the subscription model was held

back by competition - but competition has been pretty much eliminated now.

Most new Microsoft licensing schemes push in the subscription direction. For instance, the license for a recent upgrade to Microsoft Instant Messaging (L17) clearly states that Microsoft **may charge for future upgrades**, and that **you are required to purchase the upgrade when it comes out**. Your license to use the current version is terminated when an upgrade is released. While this isn't a true subscription (Microsoft still has to issue something to get paid), it's a big step in that direction, as is License 6.

On August 1st, 2002, Microsoft's volume license agreements (volume = 5 or more) changed dramatically to a "pseudo subscription". Under **License 6** (L0) companies can no longer update their Microsoft software when they please, they must run the version Microsoft dictates, and must upgrade when Microsoft says to. This is not a true subscription, but it's now just one short step to a true subscription.

Any business that didn't sign up for License 6 by the end of July 2002 is cut off completely from discounted upgrades. Should a business later decide it needs the upgrade program, **it must first purchase all new software at retail price** to get current. Higher cost Individual consumer license upgrades for home and very small business are still available until further notice, but volume license upgrades outside of License 6 are a thing of the past.

Coming as it did in the depths of an economic downturn, forcing immediate expenditure and increasing the cost of Microsoft software about 30% on average, License 6 has been unpopular (L1). Nonetheless, so many companies felt they had no choice, Microsoft's revenue increased 10% for the quarter. Extortion is a wonderful thing!

A side-effect of License 6 was to further depress the already severely depressed PC market as money had to be diverted from hardware purchase to fattening the coffers of Microsoft.

In June, 2002, surveys by "Microsoft friendly" groups were showing 40% rejecting License 6, and another 30% undecided. Many said they just didn't have the money to comply, even if they wanted to. Many were looking seriously at alternatives for the very first time (L5, L7). About 42% of Microsoft's customers scrambled to get renewed with License 5 before it was discontinued, to put off the License 6 decision for a couple of years.

In early 2003, with less than a third of corporate customers signed up for License 6, Microsoft has been making some concessions (though not big ones) to make the program easier to swallow (L15). License 6 has raised costs for 60% who participated, even double for some, and many who did not participate still don't have free cash available to even qualify. Despite the early outcry, fewer than 4% of Windows shops are actually in the process of moving entirely to competing

platforms - it's scary, and immediate costs for many may be higher than License 6.

Oh, yes, one other note: License 6 volume licensing **doesn't include any actual Windows licenses**. You still have to buy a Windows License with each and every PC you purchase (L10, L9).

License 6 is certainly not the end of the matter. If you run Windows XP, or have downloaded the Windows 2000 SP3 (which includes important security patches), **you have legally agreed to a EULA** (End User License Agreement) that allows Microsoft to **enter your computer systems**, examine them, and make changes, including **disabling software** (L13) without your knowledge or consent, and **without liability** for any damage that may result from such acts (L6).

If you doubt Microsoft is really serious about these details of license, note that they are the principal financial backer of the push for UCITA, which is intended to codify these points into state law through the Uniform Commercial Code. UCITA includes the right to enter your computer and disable software for any real or imagined violation of license or payment, and to unilaterally change terms of license after the fact, license terms which apply **even if you are not allowed to read the license before purchase**.

Eventually Palladium is to take over license enforcement. Palladium supports firm cut-off dates, so if you don't pay your subscription fees, you will not get a new Palladium key and will not be able to use Microsoft programs, **or the data you created with those programs**. In other words, you pay, or **you are out of business**.

For small business, the real revelation will be discovering just how **horrifyingly expensive** it is to run Microsoft software legally. Right now, I can count the number of small businesses I know of that are 100% legally licensed on the fingers of one hand, without using any fingers.

The cost of Microsoft software is already so high, they're offering **financing plans** (L4) to help small businesses comply with License 6. These plans will be just so easy to convert into subscriptions, once businesses become used to the monthly payments.

An ESL (End User Subscription) license for Microsoft Office XP was test marketed, and withdrawn (L11) after a year and a half because (typical for Microsoft licensing) it was too complex for users to understand, and involved dealers as well as customers. Expect ESL to be back in simplified form and with added incentives in conjunction with **Office.NET** (Microsoft never gives up on something they really want).

Even subscription is not the true objective. Microsoft's long term goal is "**Software as a Service**", through the .NET Initiative (pronounced "Dot Net"). Software won't be loaded onto your computer at all, it will run "as a service" from Microsoft .NET servers. Your business data will also reside on Microsoft .NET servers

somewhere out on the Internet. For access, you will have to be authenticated by a Microsoft Passport server.

**Bill Gates** was recently asked if Microsoft software might eventually be available only for rent through .NET, and replied "I believe in the long run things will be architected that way" (L3).

Microsoft Office .NET (N5) is far enough along that Microsoft has been doing focus group marketing studies for it (one of which escaped onto the Internet). Why will businesses sign up for Office .NET? Because it will be a lower up-front cost for each workstation, and Office .NET will offer attractive features not available in the boxed version.

So, why is Microsoft messing around with subscription software at all, instead of just going directly to .NET? That's what they wanted to do, but it's now obvious the broadband Internet access required by .NET will not be universally available in the near term.

Meanwhile, just to make sure you can't use copies of Windows you already paid for on the new computers you buy, Microsoft has forbidden computer makers from shipping computers without an operating system, under threat of "renegotiating" their OEM distribution licenses (L8).

## DOT.NET

Microsoft's .NET Initiative (pronounced "DotNet") is Microsoft's take on Web Services, and is being deeply "integrated" into every Microsoft product to assure its wide distribution and wide adoption.

.NET differs very sharply from everyone else's concept of Web Services. While everyone else concentrates on making Web Services platform agnostic, .NET is designed to **force the use of Windows to the maximum extent**. While Microsoft pays lip service to "multi-platform", they have yet to demonstrate anything of that kind.

Under .NET, Microsoft provides the key services, running on Microsoft servers and accessed over the Internet. Services not provided by Microsoft will run on ASP partner's servers. Access to any and all is cleared through Microsoft's **Passport** authentication servers.

Clients, of course, must run Microsoft Windows if they expect to be anything like fully functional. Any and all .NET services must run on Windows.NET servers - no "platform agnostic" services here, thank you.

**Java**, used as the main programming language for Web Services elsewhere, is not used with .NET, because Microsoft was a very bad boy and is **forbidden by court order** from implementing Java. To replace Java, Microsoft has had to write a new Java-like language named C# (pronounced "C Sharp").

Due to problems implementing Hailstorm services (see below), .NET is now loosely segmented into two

domains: business and consumer. The version for larger businesses has moved closer to generic Web services, with more control held by the customer (for now). The small business and consumer domain remains truer to the original "Microsoft controls all" vision.

**Software as a Service** is Microsoft's holy grail, and they are developing .NET enabled ASP versions of all their software as quickly as possible. **Microsoft Great Plains**, for instance, is on a crash program rewriting all their accounting software in C# to make it .NET enabled.

The beauty of the ASP model for Microsoft is that **all your business data**, as well as the software to access it, **resides on .NET servers**, not on your own machines. This gives them a single choke point, Passport, and tremendous leverage in raising prices (remember their revenue growth imperative). Very large organizations will have their own Passport authentication (a painful concession for Microsoft) but for everyone else, it's Microsoft servers.

Microsoft is pressuring (and sometimes bribing) Windows software publishers to quickly rewrite all their software using Microsoft's Visual Studio .NET programming tools so it will be .NET compatible. Actually, Microsoft means for many of these publishers to perish in the transition, but right now they need a show of support.

Microsoft is also heavily promoting .NET among prospective **ASP partners**, because there's no way they can provide a full range of services right from the start. Many will be intoxicated by the smell of Microsoft's money and invest in setting up .NET services to tap into that revenue stream. Their eventual fates are exemplified by **RealNames**.

**RealNames** came up with a very useful service allowing regular names of companies and agencies to be typed into Internet Explorer instead of the cryptic http://url names. This was particularly important in Asian countries because they could type in names in their native character sets and RealNames would translate and bring up the site.

In 2002, Microsoft abruptly **canceled the RealNames contract** and removed the essential link from Internet Explorer, completely ignoring loud protests from its Oriental customers. RealNames was forced to close and all its employees were without jobs (N2, N3).

Why did Microsoft kill a company with such a useful product, a company that was showing solid signs of becoming profitable? Several reasons have been given, but are not convincing, Microsoft had already hired several former RealNames employees, so I expect to see them incorporate this technology into their products without paying for it. They just wanted RealNames **too thoroughly dead to sue them**,

This same fate has befallen numerous Microsoft "partners" in the past, and awaits those that sets up

.NET services. **Either your service is not successful and you go out of business, or it is successful, Microsoft commandeers it, and you go out of business.**

All is not, however, going as smoothly as Microsoft would like. The key .NET services were bundled together into a package called **My Services** (code named **Hailstorm**). Unfortunately, My Services had to be withdrawn, because most of the big retailers and financial services refused to sign up because **they didn't trust Microsoft** (Yes, even corporate executives learn - it just takes longer). In fact, most of these organizations have signed up with the competing **Liberty Alliance** (N14), recently joined by the U.S. General Services Administration and the U.S. Department of Defense (N15).

My Services didn't die though. Within a couple of weeks Microsoft had it renamed and the marketing pitch reformulated to present it for in-house use on corporate networks. Now, they're just building it into all their products so you can't escape it (N1).

The FTC (Federal Trade Commission) has already charged Microsoft with **deception, misrepresentation** of the security and privacy of the Passport service, and with **gathering information that was off limits** (N6). Caught with its hand in the cookie jar, Microsoft responded with its usual "We didn't do it but we won't do it again" and agreed to compliance audits every two years (N11). Meanwhile, they are canceling the electronic Wallet attached to Passport and setting it up as a separate service so they can pilfer information free of the privacy and security terms of Passport (N12).

"Well, I'm just not going to be part of that", you say? If you run Windows, you probably **already are part of it**. Hot Mail, MSN, Microsoft Instant Messaging - they all require a Passport account, and they are all being expanded into key parts of **My Services**. If you are attached to the Internet, your Windows computer is already chatting with Microsoft's servers, and if your run XP or Win 2000 SP3, you have already given Microsoft permission to examine your computer and make changes as they see fit.

Windows XP is substantially integrated into .NET, and future Windows versions will be completely integrated. You won't be able to tell where your network stops and .NET starts. The licenses you need to use the software you depend on will all be authenticated through Passport and Microsoft license servers on a daily basis.

All this sounds awesome, but there are a number of factors that threaten to seriously limit the success of .NET.

Microsoft's .NET vision depends on high speed Internet access. Universal high speed access is a lot farther away than Microsoft thought when .NET was launched, so implementation will be confined much more to in-house networks than Microsoft desires. High speed access is so important to .NET, Microsoft has set up a lobbying office in Washington to encourage legislation

that would promote high speed access (N13).

Microsoft's ability to deliver viable product in a timely manner is questionable. Interactive television is an example where they have failed, yielding most of the market to competitors, many Linux-based. Two years after the .NET announcement, the Visual Studio.Net programming tools are the only .NET component substantially functional.

I expect a severe limiter to be the inability of Windows.NET to run on any but Intel based servers. The huge server farms required will be expensive, failure prone and difficult to administer. Microsoft has been trying for years to migrate HotMail servers from Unix to Windows, and the job is still incomplete. The transition crew found **Unix much easier to administer** than Windows (A15).

Security: .NET will be cracked early and often. It will bleed confidential customer information like a stuck pig (N18), and portions of it will be brought down by DDoS (Distributed Denial of Service) attacks every now and then.

Antitrust action will be brought by private parties (easier now that Microsoft stands convicted of monopoly abuse), then by the U.S. government soon after the Bush/Ashcroft administration leaves office.

Java now has a lot of momentum. It's not going to be easy to convince corporate developers to accept a Web Services platform that isn't compatible with Java, and which ties them to a single company and single class of server.

Trust, Trust, and Trust. Just about nobody trusts Microsoft any more, not even after a few drinks. Not even after smelling the money. Particularly after the screwing they got with License 6, businesses are wary of tying themselves tightly to Microsoft products. This is becoming a major and growing problem for the Redmond empire.

Microsoft has recently removed the .NET designation from a number of products and has recently described .NET as "middleware", a serious conceptual downgrading. Industry analysts are becoming impatient with Microsoft's failure to clearly define what .NET is, yet many companies are going ahead and starting major system implementations based on faith in Microsoft. The risks for these companies is substantial and the benefits uncertain (N17).

All told, I expect .NET to be widely implemented, but ever changing and never completely defined. It will not deliver fully on functionality or performance and will suffer "significant" security problems. Many customers who build their business models on .NET will suffer chronic but generally sub-fatal dysfunction. Deliberate incompatibilities with Sun's Sun One and IBM's WebSphere environments will hobble many .NET users.

Consultants, integrators and contractors will make vast

amounts of money trying to make .NET work right, but Microsoft will rake in the really big bucks.

## XML

Microsoft is enamored of XML, to the point they try hard to convey the impression that it's a Microsoft protocol, XML is an open standard for communications between systems in a Web Services environment. A subset of SGML, XML is under development by the W3C (World Wide Web Consortium) (N16).

Microsoft was involved in creation of SOAP, a remote procedure call protocol that works with XML, but David Winer of Userland was the principal architect of SOAP. Microsoft has promised complete compliance with the XML standard and did announce that the Microsoft Office file formats will be transitioned to XML.

Many people rejoiced, saying, "Microsoft is converting even Office formats to standard XML, so soon we will be free of all those problems with proprietary Office formats". This is just silly. "Extensible" fits perfectly with Microsoft's traditional approach for destroying standards: "Embrace, Extend, Exterminate".

Even a casual reading of XML specifications will show that you can define data types that require a special parser to interpret them, which could be a parser only available in a Microsoft product. Microsoft has already stated they will do this, "to protect our intellectual property".

Now it appears even that won't be necessary. Microsoft has backed off its earlier statements and now says Office 2003 will consume standard XML, but won't produce it (N8). It'll still be all proprietary formats. No prediction has been made as to when there might be an Office that saves in standard XML format.

A serious problem with this strategy is that low cost competitor StarOffice and no cost competitor OpenOffice already use standard XML as their native format and will be much easier to integrate into enterprise systems. Further, OASIS, a leading industry standards organization, is developing an XML standard for office applications, using the OpenOffice formats as the starting point (I6, I17).

If Microsoft can get business to go along with their XML schemas, and many businesses will, then business to business, and even business to consumer activity can easily be diverted through .NET services controlled by Microsoft. For that, you will pay a fee.

Intuit is currently showing the way to do this by hijacking the traffic of QuickBooks users using the built in email invoicing feature (N4). This turns out to be an ASP service for which Intuit intends to start charging a fee. Further, Intuit reserves the right to add third party advertisements to your invoices and gather all the data they want to from those invoices, right down to customer addresses and line item prices, and

use it as they please.

Expect Microsoft to be watching Intuit's work very closely, since they share the same cavalier attitude toward their customers (they tried to merge some years ago, but screwed it up).

## ACCOUNTING AND BUSINESS MANAGEMENT

Microsoft promised their business management software development "partners" they would **never compete with them**. I predicted about a year ahead of the fact that Microsoft would purchase a major Windows accounting software publisher, and predicted it would most likely be **Great Plains**, and that would be the beginning of the end for everyone else publishing accounting software for Windows.

I further predicted that a **low end accounting package** would be purchased or developed quickly to take market share away from Intuit (QuickBooks), Sage (Peachtree, Business Works, etc.), and other low range publishers. Finally, I predicted they would tack on **Point of Sale** and other more specialized modules.

Microsoft bought **Great Plains**, and six months later, Great Plains announced a low end accounting package, **Small Business Manager**. Later, in June of 2002 Microsoft purchased Point of Sale software publisher **Sales Management Systems**. The POS products will be integrated with Small Business Manager.

In announcing Small Business Manager, Microsoft made a point that it was higher priced and did not compete with QuickBooks, Peachtree and the like. This was quietly fixed some months later when a reconfigured and re-priced version of Small Business Manager was issued, moving firmly into QuickBooks territory.

Nor does all this stop with just accounting and Point of Sale. Microsoft has announced its entry into the lucrative **CRM** (Customer Relationship Management) market in direct competition with former partners. **SCM** (Supply Chain Management) is sure to follow.

For other business models, Microsoft has announced **Professional Services Automation** (B5), a software package designed for wall to wall control of professional practices, especially engineering, law and others that are "project" based. PSA integrates Microsoft Project, knowledge management, time, expense and project accounting, financial reports and analysis,

In conjunction with CRM Microsoft makes a lot of noise about concentrating on the "underserved midrange market" (\$1 million to \$1 billion in their definition) and not competing against "established enterprise partners" like SAP at the high end. Translation: "We aren't ready yet, and we still need you to lay the groundwork by converting all your enterprise clients to Windows and .NET".

When the time comes, the SAPs, People Softs and Baans will be plowed under in short order. As one CRM integrator told me, "This battle is over and Microsoft won. If you think Microsoft treats their customers badly, you have no idea how the CRM vendors treat their customers." CRM customers will go with Microsoft out of pure vengeance.

The Business Solutions division, including Microsoft Great Plains, has now been consolidated into the main Microsoft marketing machine. Given saturation of their traditional markets, Microsoft can only continue to intensify their efforts in these markets (B10). The announced revenue projections for Great Plains do not allow for surviving competitors.

Caught like deer in the headlights will be midsize customers - companies large enough to use complex systems, not large enough to be confident they can integrate those systems themselves, and fearful competitors will gain advantage over them if they don't use them. Microsoft's promise of easy top to bottom integration will be irresistible.

Microsoft's "development partners" are going through the same three stages we have seen in other markets Microsoft has invaded: **denial, desperation, bankruptcy**. They'll all get a brief mention on [www.fu??ed company.com](http://www.fu??ed company.com) (you need to fill in those ??s yourself).

Sage (Best Software in the U.S.), publisher of Peachtree, Act!, Business Works, DAC Easy, MAS 90 / 200 / 500, Platinum, Sales Logix and TeleMagic, denied any impact on its business (B11), but has already entered the desperation stage with a joint announcement with IBM that it will be a big supporter of **Linux**. Way too late - they've already forced their customers to Windows versions of their products, many kicking and screaming all the way. Peachtree, for instance, refuses to install if you have a Linux / Samba server on your network, even though it would work just fine (if you could install it). Peachtree support still refuses to talk to you if you have such a server. Sage / Best is toast.

"But," you say, "how can Microsoft displace Intuit? Everybody's using **QuickBooks**. Intuit is just too popular." Microsoft's "Triple Terminator" is the transition of Windows software to .NET, Longhorn and Palladium. Microsoft will "help" Intuit make these transitions the same way they "helped" WordPerfect and Lotus make the transition from DOS to Windows. Only Microsoft software will be able to take full advantage of the this future world - Intuit is the WordPerfect of the future.

So what does all this mean to your business? When evaluating your management and accounting options, you need to keep these points in mind:

If you stay with Windows for your business management software, you will, sooner or later, be running only Microsoft software.

Microsoft will effectively own your business, because all



your data will be in proprietary formats readable only with Microsoft products, which will be time limited (see Palladium). If you don't make your payments, you will be out of business.

Microsoft intends to move their customers to a .NET / Web Services / ASP model, particularly small business customers. Your software and business data will not reside on your computers, but on Microsoft .NET servers, with access controlled through Microsoft Passport servers.

If you bite the bullet and move to an alternative platform, you'll still have to pay for the Microsoft platform, perhaps for years, to access your historical data.

Microsoft's economic model requires rapid revenue growth. As their markets saturate, they must squeeze more money out of established customers. The whole purpose of monopoly is to increase profits through unrestricted price increases (new FTC rules have forced revelation that profits for Windows / Office are about 80%).

You won't be able to control Microsoft by threatening to go to other vendors. There won't be any, unless you are willing to bite the bullet and abandon Windows. Microsoft's goal is to make that transition so traumatic it is effectively impossible.

Even if your established competitors are as stuck with Windows as you are, new entrants into your market may not be so limited. There are mature and capable non-Windows alternatives that already cost much less, offer liberal licensing terms, do not report back to Redmond, and allow perpetual ownership of both software and data.

Think about Security

Think about The Upgrade Treadmill

By now, it should be pretty obvious you are going to have to raise your prices. 100% license compliance is going to be expensive, and you can't afford not to make your software subscription payments because that would be instant Chapter 7.

To avoid this fate, you need to start moving to alternatives now. It would have been a lot easier to do before you moved from DOS to Windows, but it can still be done if you can summon up the nerve to do it.

## SOFTWARE DEVELOPMENT

Software development may seem out of place in this business discussion, but in actual fact, only a small percentage of programmers are employed by companies that produce software for sale. The great majority work for non-software businesses developing software for **internal use**. Custom programs often yield huge labor savings compared to "canned" software.

A large amount of in-house developed software for Windows is written in Microsoft Visual Basic. While Visual Basic coders and their chosen language are roundly despised by "real programmers", they have given business a pool of relatively **low cost developers** capable of handling routine business tasks.

The Visual Basic language is very loose and forgiving. Its easy to use "point and click" structure ("point and drool" to the "real programmers") allows fairly complex programs to be created with very little understanding of programming principles.

All this changes with Visual Basic .NET, which requires much of the same planning and disciplined structure the "socially acceptable" languages have always demanded (D1, D2). This, and the fact that Microsoft is pushing very hard for conversion of everything to .NET, has left Visual Basic programmers feeling confused, disenfranchised and concerned about their futures.

The effect on businesses, as they too feel Microsoft's heavy hand pushing them to .NET, is an **increase in software development costs**. Either the existing Visual Basic programmers will have to be sent off for expensive training in the new methods (and then paid more for their new skills), or "real programmers", will need to be brought in, and they're going to want to write in "real languages" like Java or C++.

Another expense is that .NET automatically recasts all the software already developed into "legacy code" that needs to be rewritten. Microsoft provides conversion software, but the resulting .NET programs suffer a nasty performance hit and will need a lot of hand tuning to overcome that.

After that, there will be a complete rewrite for Longhorn (see above). All this is threatening to **consume the cost advantage Windows development has enjoyed over other development environments** (though some maintain this advantage has always been mythical).

These and other factors described in other sections of this document have resulted in an unexpectedly high rate of defection by software developers from Windows to Linux (D3, D4, D5). Rather than coming mostly from Unix, 52% of Linux developers previously targeted Windows. This is perhaps the most serious threat to Microsoft's continued dominance, because controlling software developers has always been their most powerful tool for destroying competition.

## WHO DO YOU CALL? - SALES AND SUPPORT

As Microsoft pushes its market toward the enterprise data center, they increasingly encounter IS (Information Systems) departments that have very high expectations of vendor support - expectations born of a long association with **IBM**. To meet these expectations, Microsoft established MCS (Microsoft Consulting Services).

Very quickly, **MCS** became as much a problem as a

solution. Microsoft's "channel partners", VARs (Value Added Resellers) and integrators, long established in the enterprise market, were finding themselves competing directly with MCS for business from large accounts. These channel partners expressed their concern very clearly.

Now, here's Microsoft's problem - .NET. Microsoft's Web services are still pretty much "vaporware" (often impolitely called .NOT), and .NET has serious competition. Sun Microsystems' Sun One and most especially IBM's WebSphere are more mature, and are backed by companies long established in the data center space. WebSphere is backed by IBM Global Services, a worldwide operation with an estimated 2002 revenue of nearly \$40 Billion (after absorbing PWC Consulting).

It would take Microsoft decades to build an organization competitive with IBM Global Services, so they **must** depend on their channel partners - not only to push .NET, but to convince customers to delay Web services implementations based on competing products available now. An upset channel could easily switch to non-Microsoft products.

"The channel" has indeed always been Microsoft's secret weapon - a huge number of VARs, resellers, consultants and integrators of all sizes and stripes promoting Microsoft products to businesses, and Microsoft has always treated these "channel partners" very well.

Now, in recognition of their intense need for the channel to do the heavy lifting, not just for .NET, but for business management and accounting, CRM (Customer Relationship Management) and SCM (Supply Chain Management), Microsoft has had to revamp their channel partnership program (S2).

A major feature of the new program recasts MCS from a profit center to a "satisfaction center" (without, of course, giving up profit). It has been decreed that MCS will always seek to be a subcontractor behind a partner rather than the prime contractor. MCS will also sell support to the partners in substantially discounted lumps of \$68,000/year for up to 400 hours (\$170/hr) and \$20,000 lumps for SMB (Small, Medium Business) partners.

This new program has a number of strategic advantages for Microsoft.

It eases the concerns of all-important enterprise partners, preserving the channel.

It places no limit on the growth of MCS.

It gets customers accustomed to dealing with Microsoft people directly.

It gets partners accustomed to Microsoft being directly involved in their deals.

Yes, Microsoft still covets its channel partners'

revenues, and grew its direct sales team by 47% in 2002, and account management staff by 15% (S4, S1), and there's still an imperative to grow MCS to satisfy their largest customers, but they must do all this very delicately, lest upset partners bolt to IBM/Linux and Sun alternatives.

The big question is, can they pull it off? Partners will still be running up against MCS to some extent. Revenue for upgrades is a major item, and is already migrating to Microsoft direct, and that will accelerate as Microsoft moves licensing to a subscription basis. I expect established channel partners to remain concerned enough to build significant capability with alternative products, but not concerned enough to jump off the Microsoft ship any time soon.

**What does this mean to the small business manager?** It means you will be dealing with your local service providers, not directly with Microsoft, for some time yet. Yes, "A Microsoft partner is a victim they haven't gotten to yet", but the channel partners still have a while to live. Microsoft just isn't big enough to do without them, and may well never be that big.

## HOME AND ENTERTAINMENT

With its business monopolies saturated, the home PC market stagnant, and heavy resistance in the corporate data center, Microsoft sees the huge **home entertainment market** as its next big revenue generator. Here they expect they can leverage their PC monopoly effectively. As with any market Microsoft enters, "we only want our fair share - all of it".

Microsoft's expressed intent is for **all family entertainment** to be delivered through Microsoft controlled channels - both hardware and software, and to charge both the content distributor and the content consumer for the privilege.

The Xbox game console introduces the concept of Microsoft manufactured and controlled hardware, but is just the first (multi-billion dollar) step. Having seized complete control of hardware design from Intel, Microsoft is setting the stage for taking over the home market entirely - disk, box and content.

Microsoft's next step is piloting now - **XBox Live (H3)**, an on-line gaming system. Unique to Microsoft's system is that **Microsoft owns all the servers**, thus Microsoft collects all the money and all the personal information about the players. **Sony** encourages third party businesses to set up and run gaming servers.

Sony is providing an adapter kit supporting both modem and broadband (H4) connections, while Microsoft supports only broadband. Remember how important broadband is to .NET and you see why they discourage modem use. Consider also that **Bill Gates** and Microsoft co-founder **Paul Allen** have invested heavily in cable companies.

Next up is the **XBox Home Gateway** through which all

family entertainment will filter. It will, for instance, contain the TV schedule (updated daily from Microsoft.NET) and will select the channels for your new digital TV. This will enable Microsoft to **sell "positioning" to TV shows** the same way they sold icons on the desktop to Internet services. Any content critical of Microsoft will be pretty hard to find.

Following Home Gateway will be **XBox 2**, which has been described as much more like a full function PC. This will be the beginning of the end for makers of PCs for the home market. They will find it impossible to build machines compatible with Microsoft's content controls.

The Home Gateway will extend, through **Universal Plug and Play** (already a default setting in Windows XP (and already a severe security problem (H1))) to control your home appliances. At least it'll give you some nifty new excuses: "A hacker broke into the microwave and ruined the turkey", "The refrigerator got a virus and that's why the food spoiled and everyone got sick".

Of course, Microsoft will be selling plenty of information about you and your family to advertisers, market researchers and probably Homeland Security. They've already contracted for profiling software for their current TV set top boxes (H2).

The Home Gateway, coupled with Digital Rights Management software, and eventually Palladium is half the "Trustworthy Computing" picture Microsoft is promoting to the MPAA (Motion Picture Association of America) and RIAA (Recording Industry Association of America) as an unbreakable distribution method.

The other half of the picture is the servers used to feed digitized content (motion pictures, music, etc.) to PCs and other devices equipped with Microsoft DRM technology (H10). A more advanced DRM Server (H5) is planned for release, and will replace the current Media Rights Manager server. The servers are used by content providers .

To support all this, the Microsoft EULA (End user license agreement) you have "signed" (by the act of using Windows XP or by downloading recent Service Packs and security fixes) specifically states that Microsoft has the **right to inspect software on your PC and to change or disable that software as they wish, without notice to you, and without liability to Microsoft**, to protect copyrights (including their own). This means Microsoft has complete administrative rights to your PC - home or business.

Every household with children or teenagers will have its computers infected with Microsoft's DRM schemes within a few months. Kids have always received everything free from their parents, and just can't understand that "free stuff" from the outside world must be regarded with suspicion. Microsoft is distributing their DRM system with free "preview" CDs and by similar means (H9, H12).

Of course, advanced DRM requires hardware support at your end, but hardware manufacturers are all perfectly

willing to go along ( H11). They see no downside whatever to your need to buy new video and sound cards, or even new PCs to support Microsoft's schemes.

All this has caused considerable unrest among rights and privacy advocates, because it **tramples a number of legal rights**, such as the "fair use" doctrine, post sale usage of purchased content, privacy, the security of your PC and control of the content you can view and the software you can run - but with lawmakers already bought and paid for, there's not a lot that is likely to get done about it.

So Microsoft's push to control home entertainment seems unstoppable, yet right in the midst of it, they received a **stunning setback**. It was announced, on September 9th, 2002, that Movielink (H8), a joint venture of Metro-Goldwyn-Mayer, Paramount Pictures, Sony Pictures Entertainment, Universal Studios and Warner Brothers would offer movie downloads starting in Q4 of 2002. **IBM**, not Microsoft, provides hosting, system operation, rights management and network management.

This deal is so big the Department of Justice jumped in before it was even announced. It's now obvious these movie giants never intended to consider Microsoft's delivery plan. Why? Looks like it's that pesky "trust" thing again. You just don't hand a company with Microsoft's ambitions control of your air supply (I told you these guys were smarter than the software industry).

The music industry, on the other hand, is a different matter - all the greed with none of the brains. I expect **Hilary Rosen's RIAA** to go hook line and sinker for Microsoft's promises. If they had two synapses to rub together they would have seen the power of Napster and turned it to their own advantage, but now, they'll be completely dependent on Microsoft for delivery, and that's **going to cost them plenty** (H12).

## SECURITY

Security is one of Microsoft's most serious weaknesses. Worm, virus invasions, trojans and now root kits, Web page defacement, credit card theft, data theft, espionage and destruction. These are all major features of Windows systems, costing business tens of billions of dollars per year worldwide.

The vulnerability issue has become so serious a public relations problem it threatens Microsoft's expansion into the enterprise datacenter and acceptance of .NET. Even long time ally Gartner Group has recommended dumping Microsoft's Web software (X34). This, not your costs, is finally getting Microsoft's attention.

Microsoft claims attacks focus on Windows because Windows is so popular, True, but it's also because Windows is a uniquely soft target. While no system is entirely secure, most at least take some skill to penetrate. Windows provides easy success for neophyte crackers and entertainment for thousands of unskilled

"**script kiddies**". Don't believe me, though, listen to Microsoft.

"... **Our products just aren't engineered for security.**" (X0) - Brian Valentine - Microsoft senior vice president for Windows development. Another Microsoft executive recently explained they never paid attention to security "**Because customers wouldn't pay for it until recently**" (X23). In other words, customers wouldn't pay **extra** for something they expected as part of the product.

**Windows XP**, heavily advertised as "The Most Secure Windows Ever" needed a **major security patch within weeks**, XP's password security is easily bypassed by even normal users (X58). Computer sophisticates correctly point out that no system is secure if you have physical access, but it just shouldn't be that easy.

The skill level needed to write a successful Windows worm or virus is absurdly low. The people who launched the famous **Love Bug** (estimated \$8 Billion in damage and eradication costs), had only a few weeks of computer training.

Even commercial products take advantage of Windows' weaknesses. Anyone can install a keyboard logger on someone else's Windows PC and have it email all activity (including network passwords) to the perpetrator's mailbox (X10, X9).

**Microsoft's .NET Initiative** will magnify the current problems, since business systems participating in .NET have to be permanently attached to the Internet, with key services coming from Microsoft servers of questionable security (X28, 57). Key features are "single sign-on" and "tight integration", so when security is broken at one point, it's broken everywhere.

Now that Microsoft's customers are desperate enough to pay extra, Microsoft has opened a **Microsoft Security Business Unit**, and is exploring ways to charge customers for add-on security products. Of course, Microsoft's existing security product, ISA Server (Internet Security and Acceleration Server) requires occasional security patches (X69).

Microsoft's much hyped firewall software for Windows XP and Windows 2003 isn't doing too well either, and Microsoft itself advises you to acquire some other publisher's software to plug the holes (X78).

What makes Microsoft's environment so uniquely vulnerable to invasion, subversion and security leaks, and why are the vulnerabilities not getting fixed?

**Single User Roots** - Microsoft's software was originally developed to run on a single user computer not attached to a network of any kind, never mind the Internet. Key integration features, such as OLE (renamed COMS, then DCOM (distributed COMS), then Active X) were originally single user code requiring no security.

You can't just tack security on to a fundamentally

insecure structure and have it be effective. Unix (and by extension Linux) and mainframes were designed to be multi-user and networked with prudent built-in security from the very first.

**Ease of Use** - Effective security is always at odds with convenience. Microsoft's main marketing pitch is "ease of use". and "integrated environment". For this reason, what security features there are are often turned off by default.

**Tight Integration** between products allows a rogue process to move freely. Tight integration through secret programming interfaces locks competitors out and is a critical "ease of use" factor as Microsoft products become more complex, so it's here to stay.

**Automation Features** - the very features that make it possible for non-programmers to automate business processes and improve ease of use are used by worms and viruses for their own automation needs. Because every Windows computer has these features, worm and virus writers can take full advantage of them.

**Intentional Vulnerability** - Microsoft **demands access** to your computers and network over the Internet, without your consent and without your knowledge (it's in the license agreement for Windows XP and Windows 2000 SP3). To expect a computer system outfitted with these "features" to be in any way secure is purest fantasy.

**Deliberate "Back Doors"** - It is strongly suspected that Microsoft has provided government agencies with keys and codes allowing secret entry into Windows systems. This is a major reason why China and the German Bundeswehr are dumping Windows. Further, Microsoft developers have hidden whole games within the code of Microsoft Office. It's absurd to think they haven't put in a few convenient entry points.

Many think the sudden surrender of the Bush Department of Justice to Microsoft, after the antitrust case was decisively won and upheld on appeal, was in return for Microsoft inserting access points into Windows - access points that are probably redundant, but asked for as a result of the now famous lack of communication among U.S. security agencies.

**Uniformity** - (X59) - In agriculture, the "Uniform Windows Experience" would be called a "monoculture". Monoculture crops, all from exactly the same genetic material, maximize ease of marketing, but are vulnerable to complete destruction by pests tuned to that exact crop. Monoculture agriculture makes heavy (and sometimes ineffective) use of expensive pesticides. Monoculture computing makes heavy (and often very ineffective) use of expensive antivirus and invasion detection software.

**A Corporate Culture** of aggressive competitiveness both within the company and without has different groups working on similar projects. They rush their development and do not communicate with others, knowing that one group will survive and the other will

not. Adding new features takes precedence over all other work, because that's what wins political battles.

**Poorly Trained Administrators** - Microsoft's constant pitch to business is that with their products you don't need highly trained and experienced administrators (color them expensive), it's all "point and click". Unfortunately, system security is rocket science, so most Microsoft shops do not have the skills to implement effective security.

**Inexperienced Developers** straight out of college are Microsoft's choice, so their employees are vulnerable to indoctrination in Microsoft's cult-like corporate culture (the uniformity of Microsofties in appearance and attitude is startling). These developers have no real experience in business methods and little interest in or understanding of security issues.

### What is Microsoft Doing About This

**Denial** has always been Microsoft's "remedy of choice", blaming "dumb users", "criminal hackers", and "poor administration" for security problems. Apparently Microsoft itself employs plenty of dumb users and poor administrators, because they've had to disconnect their internal systems for major cleanups with every big worm attack, and they got slammed pretty hard by the **Slammer worm** in JAN 2003 (X51, X48).

**FUD** (Fear, Uncertainty, Doubt) is Microsoft's second line of defense. A large number of articles are now appearing in news and magazines purporting to show that all environments have about the same level of vulnerability (X27), or that Linux has more bugs than Windows. They are typically by "Writers for Hire" (X26), and are simple rewordings of the Microsoft line.

Aberdeen, a market research firm numbering Microsoft among its most important clients, published a "report" with a catchy title claiming that Linux has more security flaws than Windows (X49). Cert, the security clearing house upon whose numbers the "report" was based, immediately declared Aberdeen's interpretation **meaningless and invalid**, but this hasn't stopped the "report" from being widely quoted.

When a security problem is found Microsoft may do nothing until it becomes a PR problem (X79, X21). Then a patch is issued, which may work, but may cause problems (X25, X56). If it isn't easily fixed, or the fix would violate Microsoft's marketing plans, they deny it's a problem (X7, X3, X4, X13).

### Trustworthy Computing

The normal measures having proved inadequate to quiet customer concerns, Bill Gates, in January 2002, issued his famous internal memo (obviously designed to be leaked) proclaiming Trustworthy Computing, and declaring security Microsoft's number one development concern (X74).

Microsoft also created a new vulnerability category, "Important", to reduce the number of "Critical"

warnings. Since then, public pressure has forced them to upgrade several "Important" warnings to "Critical", and we're still getting a "Critical" about every two weeks. (X24, X29).

Microsoft announced they were halting development work for the month of February for intensive code review and employee security training. Serious security problems continued to be discovered however, and of the first 15 found in 2002, only one was found by Microsoft. Patches have continue to roll out in an incessant stream, (X5, X8, X11, X12, X13, X14, X15, X18, X19, X22, X29, X31, X32, X43, X54, X55, X67, X69, X71, X73, X82, X87, X89, X90, X91, X92) - and many, many more, but you get the picture.

On the first anniversary of "Trustworthy Computing", Bill Gates boasted of spending over \$100 million reviewing code, and an ad campaign was kicked off promoting Microsoft security (X47). The effectiveness of this expenditure immediately came into question. Not only have security problems continue to be announced apace, some truly extraordinary events followed Bill Gates' anniversary speech.

The **Slammer Worm** infected every vulnerable computer in the world **within 10 minutes** of launch, exploiting a port address carelessly left open by the default installations of Microsoft SQL Server.

In the U.S. most of the damage was contained in the early morning hours, but many found they could not use bank **ATM machines** (X52) or **airline reservation systems** for the rest of the day. Unlike Nimda and Code Red, Slammer did not become a chronic problem, simply because it was so severe. If you had an infected machine, you had no choice but to just turn it off (X46). Once again, it could have been much worse, but the worm's designer didn't optimize for damage (this worm was only 376 bytes in size).

Countries with more centralized Internet structures than the U.S. has, like South Korea, lost Internet access entirely for some hours. In the U.S., access was just deathly slow (I experienced it first hand at about 1:00AM). Had Slammer been designed to do damage, Internet access might have been unavailable for weeks, and the cost astronomical.

It's just coincidence, but the **Department of Homeland Security** moved its servers from **Windows** to **Linux** on the very day of the Slammer attack (X45). Note also the great improvement in uptime since then.

In March, 2003, The **U.S. Army** was hit by an attack exploiting a vulnerability that had not been previously announced by security researchers (X67). Most attacks exploit well known problems. Microsoft rushed out a patch, but it killed some Windows 2000 servers (X68).

Windows NT, XP, 2000 and the "Trustworthy" Windows 2003 Server were all found vulnerable to an exploit giving system level control to an invader (X88).

Trustworthy Computing may eventually produce a

"better moustrap", but "better mice" are here already. **Root kits**, a far more sophisticated attack than traditional methods, are now available for Windows (X66). A properly written root kit can be nearly impossible to detect, but a few have now been found in the wild (mainly due to programming errors).

In May 2003, Microsoft's **Passport security system** was shown vulnerable to a simple to use, yet extremely severe security exploit (X75). Passport is where Microsoft wants you to store all your credit card numbers and stuff for "single sign on" access to financial transactions. All that was needed to seize control of a Passport account and all it contained was to know the owner's email address. A month later, another exploit was found (X86) that would allow hijacking accounts.

This could go a long way towards explaining how spammers have been hijacking thousands of HotMail accounts to broadcast unsolicited email. That such a major yet easy to find flaw escaped Trustworthy Computing code review casts doubt on the whole process. Gartner Group has advised their corporate clients to shut off any Passport services they may be using (X77).

HotMail itself has been riddled with problems, such as the "spammer hole" which was unfixed for months (X83). Microsoft finally "fixed" it by limiting HotMail users to 100 emails a day.

Due to an agreement with the Federal Trade Commission resolving a previous Passport security issue, Microsoft could be subject to fines of up to \$2.2 Trillion (FTC figure) (X76). Any other company could expect severe punishment, but this is Microsoft, so everyone expects them to get off without paying a dime.

In May 2003, **Internet Explorer** was shown to **shut down instantly**, if a user views a Web site into which a simple one-line bit of HTML code is embedded (X73). Mozilla, Opera and other competing browsers ignore the bad code.

In August 2003, the Blaster worm went wild, crashing computers and disabling entire networks (X90). It then launched a denial of service attack against Microsoft's Windows Update server, but, fortunately for Microsoft, the worm writers used the wrong address so the attack was easily deflected.

A few days after Blaster launched, the W32/SoBig.F worm was released, displacing W32/Klez.H as the most virulent worm ever (X91), and Microsoft announced two more critical flaws in Internet Explorer (X92).

While "Trustworthy Computing" isn't doing well on the security front, customers have started complaining that the initiative is being used **primarily as a marketing tool** to force them into unwanted upgrades to newer, "more secure" versions of Microsoft products (X33). This is certainly the case with the release of Windows 2003.

#### **Patches**

Why are so many attacks able to exploit flaws long known and for which patches are available? For one thing, Microsoft has worked hard to convince people Windows will save money because it doesn't require skilled administrators, so few businesses have them, and security patches aren't tracked and don't get applied.

The **Code Red** and **Nimda** server worms came out long after the patches, but a year later servers all over the world still transmit these infections. When informed of an infection, server owners become very angry at the messenger, but don't fix the problem - most have no idea how.

Companies that do have skilled administrators are swamped by the sheer number of patches, and difficult patching procedures, some taking hours to apply. Microsoft itself has admitted the problem at TechEd (3 June 2003) and promised two new patching programs, one for Windows and one for applications, but those won't be available until the end of 2003 at the earliest (X81).

Many are reluctant to install patches at all because they **often break something else**, or make the system unstable or nonfunctional (X53, X68, X72, X80, X85). Sometimes the patches simply don't fix the problem anyway. Sometimes they are used by Microsoft to force new license terms customers don't want, as with Windows 2000 SP3, and to install new features that may disrupt operation.

For desktop computers directly attached to the Internet, the situation is infinitely worse. **Most will never be patched**, and make convenient tools for **DDoS** (Distributed Denial Of Service) attacks and as entry points into company networks. For this reason, Microsoft is trying to force feed updates **using Windows Update** every time you attach to the Internet. Anything wanted by Homeland Security will ride along, of course.

#### **Information Gathering**

A real concern for many users is **just exactly what information is being sent to Microsoft** during a Windows Update session. Originally, Windows update sent no information at all, but this has changed. A German organization has recently examined the data stream with special tools and found a lot of information about your computer is now going to Redmond, including the Windows product ID and a listing of its hardware configuration (X60).

The usefulness to Microsoft's licensing efforts of a listing of software on your computer is obvious. It's not currently included in the Windows Update transmission **but that capability is already in the product** (X70), to be turned on any time Microsoft wants it.

The on-line registration wizard for Windows95 was suspected of sending such a list, but Microsoft denied it. The data stream seemed to show no list, but later analysis revealed numeric codes were assigned to about

100 software packages, and the first 10 found on your computer were included in the header block of the transmission (X61).

Unfortunately, for Windows XP users, the problem **goes far beyond Windows Update**. While XP will work without an Internet connection, if it can get one, things change. The Application Layer Gateway process configures your computer so it can be controlled by Microsoft's computers. XP has at least 16 processes that communicate with Microsoft and at least 11 processes that update software from Microsoft's servers, all without your knowledge or permission.

A lot of information is sent from your XP computer to Microsoft. For instance, if you watch that DVD of Debbie Does Dalas using Microsoft Media Player, that information is sent to Microsoft, as is various information about user skills and computer usage.

All this can be turned off by rather tedious processes, but as soon as you install a security patch or other update from Microsoft, it'll probably all be turned back on. If you attempt to prevent XP from communicating with Microsoft using Zone Alarm or a similar software firewall, XP will disable features of your computer. If you are not paranoid about this yet, you should read all the details (X84).

## Conclusions

All in all, security experts aren't real impressed with Microsoft's "Trustworthy Computing" efforts (X50). In fact, some programming houses have adopted the phrase "Trustworthy Computing" to mean bugs, as in "Make sure there's no Trustworthy Computing in that code". Some experts feel Windows is broken so severely at such a deep level it simply cannot be repaired (X1, X2).

If your business has data it must keep secure and protected, such as medical records, legal files, or credit card numbers, or your systems must be always available, or if you have something to hide, running Windows leaves you vulnerable to data loss, theft, and possibly to arrest or legal action for failure of due diligence.

If you are a lawyer, why bother with the criminals - charging the victims with negligence is going to be much easier (I am not a lawyer, so this does not constitute legal advice). If you have criminal intent, or antisocial tendencies, I suppose you already know what to do and are probably already doing it.

Next Installment will talk about Microsoft's legal woes and the competitive landscape.

(A full list of citations will appear with the last installment)


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

<http://www.aaxnet.com/editor/edit029.html>

# GIMP: Make Your Own Brushes

Author: Eric Lamarque <[eric\\_lamarque@yahoo.fr](mailto:eric_lamarque@yahoo.fr)>

## INTRODUCTION



In this article I will present the various brushes that GIMP offers and show how to create brushes.


First we will look in detail at the brushes dialog, "Brush selection", and the different brushes it

contains. After that I will show you how to create your own brushes.

The version of Gimp used for this article is 1.2.3.

## 1. THE AVAILABLE BRUSHES

First, if you have not done it yet, open the dialog "Brush selection" (from Menu/File/Dialogs/Brushes). The following window should appear:



1. Size of the brush: here 13x13
2. Name of the brush: Circle Fuzzy (13)
3. Current brush
4. This small '+' indicates that the brush is larger than shown; press and hold the left mouse button on these brushes to see them in their real size.

5. This small red triangle indicates a, so called, **pipe** brush: this brush contains several images and not only the one shown; click with the left mouse button on these brushes to see all the images.
6. Value of spacing. It is the size of the brush in percentage (here 25%) which GIMP uses before the next copy of the image is used. (No panic: you will see how that works).

Now let us see the various types of brushes. For all the exercises in this part, open a new image RGB with a white background.

Reset the colors of foreground (FG) and background (BG) to black/white. Select the brush "Circle Fuzzy (13)" and draw a line with the paintbrush tool. Change the color of foreground to blue and draw a second line.

Impressing!!

It's a **simple brush**. The brush in this case is a grayscale image used by Gimp as an inverted alpha channel. The white corresponds to the transparent color and the levels of gray are used as value for the foreground color. The file associated to the brush has the extension ".gbr".





Set again the colors of foreground and background to black/white. Select the brush "Vine" from the "Brush Selection" dialog. Change the foreground color to blue and draw something.

You note: you wanted vine in blue but it remained green. "Vine" is a **RGB pipe** or **animated brush**: it is not possible to change the color in this type of brush. On the other hand, the image evolved/moved while you draw. The associated file type for this brush has the extension ".gih".

Reset the colors of foreground (FG) and background (BG) to black/white. Select the brush "PencilSketch" and draw. Change the color of foreground into blue and draw a second line.



Ah finally! This starts to resemble something! "PencilSketch" is a **pipe "alpha"**. In addition to drawing with the color of foreground, the brush changes during the drawing. The file type associated with this brush has also the extension ".gih".



• Change "Spacing": Reset the colors of foreground (FG) and background (BG) to black/white. Select the brush "Circle(13)", set spacing to 20% and draw a line.


• Set spacing to 100% and draw a second line.

• Set spacing to 200% and draw a third line. This experiment shows the influence of spacing on the brushes. For decorative brushes ("Guitar "or" Pepper"), one will generally choose a spacing of 100%.

## 2. CREATION OF A SIMPLE BRUSH

Reset the colors of foreground (FG) and background (BG) to black/white. Open a new image (File->New) with image **type gray** and fill type **background**.

**Trick:** Do not hesitate to create an image larger than the size of the desired brush: it is easier to work on a big image and then reduce than to make the brush in real size. For the example, the image with a size of 256x256 and will be reduced to 64x64.



For the example, I drew a boat.

**Trick:** If you want to obtain a brush with soft contours, use the filters "Blur" from the right click menu.




Your master piece is ready. Here an outline of my example already reduced to 64x64.


If you think of modifying/reusing the elements of your image later on, it is the moment to save now as .xcf or **never come back**: the creation of the brush is an operation with a single direction. Once the file ".gbr" is generated, one loses the original elements of the work in progress (path, channels...).

Now to finally create the brush do this:

- if your image contains more than one layer, transform it so that it has only one layer and flatten the image (layers->flatten).
- if it need be, scale the image to the final size of the brush. In the example, the image is reduced to 64x64 (right click image->Image->Scale Image).
- save your image with the extension ".gbr" in the directory ~/.gimp-1.2/brushes/. In the example, the file is called "bateau.gbr".



- give your brush a name and choose spacing (percentage). These parameters will appear in the dialog "Brush Selection". Here, the brush will be named "bateau" (French for boat).



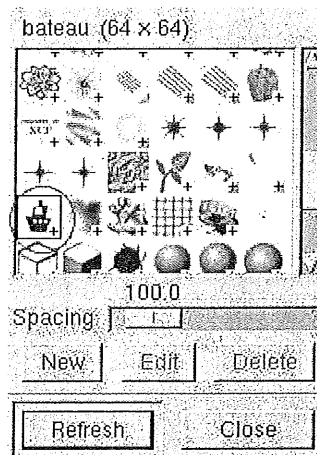
And ready is the magnificent



decorative brush. But it does not appear in the dialog "Brush Selection".

Press the refresh button and it becomes available.

You can now use your brush.

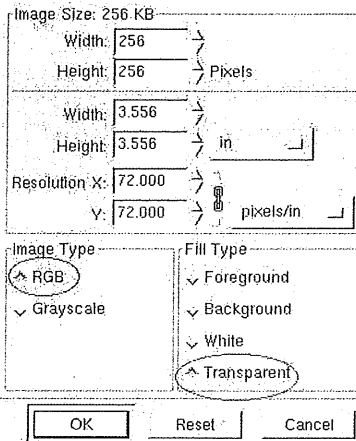




### 3. CREATION OF A RGB PIPE

A pipe - or animated brush - is a brush which contains several images. The image of the brush varies at the same time as one draws: it is the experiment we did with the brush "Vine" in the first part.

In GIMP, the creation of a brush of this type is an image with several layers.



**NOTE:** Only the brush "pipe" can have colors (RGB). If you want a fixed brush color (like "Pepper"), you will have to create a pipe with only one layer.

This is what you do: create a new image **RGB** with "fill type" **transparent**.


**Hint:** Do not hesitate to create an image larger than the size of the desired brush: it is easier to work on a big image and then to reduce than to make the brush in the final size.

Remove the layer which is by default there. It is called "Background".

And here we are, using our imagination, we just have to create the brush of our dreams.. For that, create a layer with a "transparent fill type" for each image of the future brush.

For the example, I used my image of the boat with 3 different colors.

**Important:** All the copies must contain an alpha channel. The simplest way to do this is to create the copies with a transparent background ("Fill Type=transparent").



Your master piece is ready. Here a screenshot of the layers from the example: a boat with 3 colors.


If you think of wanting to modify/re-using

the elements of your image later on, then save it now as .xcf : the creation of the brush is an operation with single direction. Once the file "gih" is generated, one loses the original elements of the work in progress (layer, path, channels...).

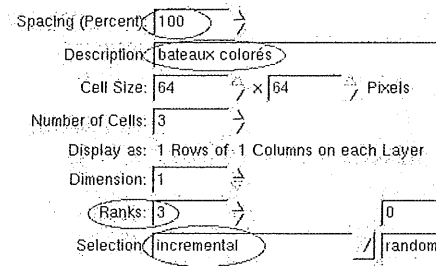
Now to finally create the brush do this:

- if need be, scale the image to the final size of the brush e.g 64x64 (right click->Image->Scale Image).
- save your image with the extension ".gih" in the

directory ~/gimp-1.2/brushes/. In the example, the file is called "bateau.gih".



The dialog to create a pipe has a lot of possibilities. For clearness of the article, I will just focus on the simple case:

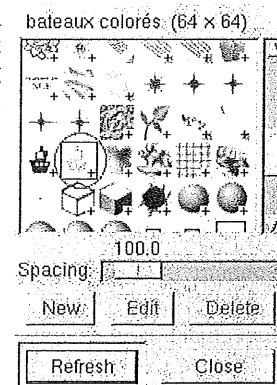


1. Enter the spacing in percentage
2. The name of the brush as it should appear in the "selection of brushes" dialog.

3. Ranks: this parameter must correspond to the number of layers present in the brush (4 layers in the example).
4. Selection: it is the order in which the layer(copies) of the brush are used. In the example, "incremental" will produce the image "Deep Blue" then "Blue" then "Blue blade" then "Deep Blue"... The other modes which are easy to use are "angular" and "random".

It is done: now you have a new animated brush. But it does not appear in the dialog "selection of brushes" until you press the refresh button.

You can now use your brush.



### 4. CREATION OF AN 'ALPHA' PIPE

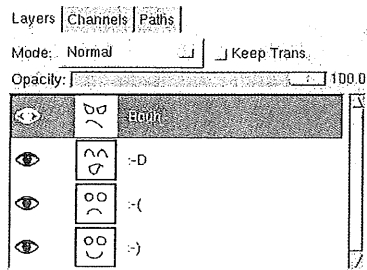
A "alpha" pipe is an animated brush which contains several images of the type grayscale. The brush "PencilSketch" is of this type.

In GIMP, we need to create an image with multiple layers and image type "Grayscale".

Reset the colors of foreground (FG) and background (BG) to black/white. Create a new image of type "Grayscale" and set fill type to "Background".

Remove the layer which is by default there. It is called "Background".

And here we are, using our imagination. I drew for example a simplified face. **IMPORTANT:** No layer must have any transparent pixels. The easiest way to achieve this is to create all layers with the fill type "background".




The master piece is ready: four different expressions of a face.

If you think of modifying/reusing the elements of your image later on, it is the moment to save now as .xcf or **never**


**come back** : the creation of the brush is an operation with a single direction. Once the file ".gih" is generated, one loses the original elements of the work in progress (path, channels...).

Now to finally create the brush do this:

- If need be, scale the image to the final size of the brush. In the example, the image is reduced to 64x64 (right click ->Image->Scale Image).
- Save your image with the extension ".gih" in the directory ~/.gimp-1.2/brushes/. In the example, the file is called "faces.gih".



The dialog to create a pipe has a lot of possibilities. For clearness of the article, I will again focus on the simple case:



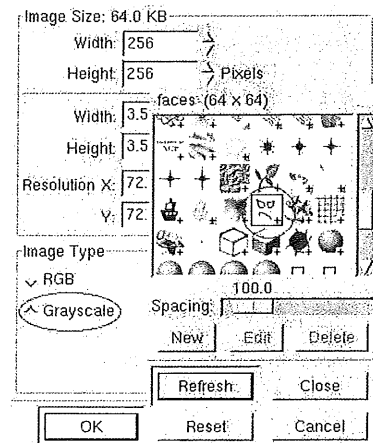

1. Enter the spacing in percentage
2. The name of the brush as it should appear in the "selection of brushes" dialog.
3. Ranks: this parameter must

correspond to the number of layers present in the brush (3 layers in the example).

4. Selection: it is the order in which the layer(copies) of

the brush are used.

It is done: now you have a new animated brush. But it does not



appear in the dialog "selection of brushes" until you press the refresh button.

You can now use your brush.

## 5. IMPORTING OF BRUSHES FROM PAINT SHOP PRO

There are a lot of brushes available for Paint Shop Pro (extension ".tub" or ".psp") and it may be that you find one which correspond to your needs.

1. Check first of all that the conditions/license to use the brush are compatible with your project.
2. Simply open it in GIMP; if this fails, then the format is unknown because the file was created with a too recent version of Paint Shop Pro. No chance.
3. When brush is open in GIMP; save it with the extension ".gih". The majority of the parameters are already correct just set the name of the brush and the spacing.
4. Move the brush image to the directory ~/.gimp-1.2/brushes/. Press the refresh button in the dialog "selection of brushes".


Let's try and example from [graphicssoft.about.com](http://graphicssoft.about.com), the site of Sue Chastain.

This site has the advantage of providing good information for each brush: note it down for later use.


Brush "Translucent Colored Orb Tubes" (file sc-colororb.tub):

- "Total Cells 8": the pipe contains 8 images
- "Cells Across 4, Cells Down 2": lines in 4 columns and 2 rows
- "Step Size 150": spacing with 150 pixels

Open the file in GIMP. Even if the extension ".tub" is not mentioned explicitly in the type of files, the format Paint Shop Pro ("PSP") is supported.



You should now have an image containing eight bubbles colored on a transparent background.



Save the image under the name "sc-colororb.gih".

Spacing (Percent): 100 →

Description: SC Colored Orb

Cell Size: 146 → x 150 → Pixels

Number of Cells: 8 →

Display as: 2 Rows of 4 Columns on each Layer

Dimension: 1 →

Ranks: 8 →

Selection: angular → / random

I marked in red circles on the left side the fields you need to modify:

1. Spacing: 150 pixels. With the bubbles having a size of 146x150,

we can assume that it will be 100% (a full bubble).


2. The name of the brush: SC (the initials of the author - rather widespread convention) and some meaningful name.

Notice that GIMP automatically filled the number of cells (8) and the way of finding them in the image ("2 Rows of 4 Columns": 2 lines out of 4 columns).

The number of rows and the selection are also converted (8 rows - angular selection).

Save the file with these settings and copy sc-colororb.gih to ~/.gimp-1.2/brushes/.

In the dialog "Selection of brushes", use the button "refresh" and your new brush is accessible.



You can now use your brush.

## 6. CONCLUSION

Now, you know almost everything about brushes. But for what can that be useful? Here some examples:

- To convert the characters of the dingbats fonts into decorative brushes ( [Playing around with Dingbats and The Gimp](#) by Katja Socher)
- Retouch/improve photos ( [Digital Hair Mania](#) by Russell Brown)
- To draw along a path ( [Stroking along path](#) by Gautam N. Lad)

One can also create brushes to change photographs, to add texture to parts of an image and probably plenty of other uses.

Now, with your brushes!

## REFERENCES/LINKS

**Note:** The tutorials ones on [mmmaybe.gimp.org](http://mmmaybe.gimp.org) will change URL at the time of the redesign of the [www.gimp.org](http://www.gimp.org) site: [mmmaybe.gimp.org](http://mmmaybe.gimp.org) will become [www.gimp.org](http://www.gimp.org).

- GIMP User Manuel - Creating A New Brush (Karin Kylander - Olof S. Kylander) <http://manual.gimp.org/manual/GUM/brush2.htm>
- Custom Brushes Tutorial (Gautam N. Lad) <http://mmmaybe.gimp.org/tutorials/CustomBrushes/>
- Image Pipes (Adrian Likins) <http://mmmaybe.gimp.org/tutorials/ImagePipes/>
- Image Pipes 2 (Sam Jones) <http://mmmaybe.gimp.org/tutorials/ImagePipes2/>
- [Graphics Software](#), the site of Sue Chastain who gave me permission to use her colored bubbles.

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

<http://www.linuxfocus.org/English/November2003/article314.shtml>

# Linux Radio Station

Author: Phil Hughes <[phil@ssc.com](mailto:phil@ssc.com)>

## INTRODUCTION

In late September I wrote an article about radio station automation for the *Linux Journal* web site. You can find the article [The comments I received](#) indicates that there is interest. What I would like to do here is to see if there are some people who would actually be interested in working on such a project. I have established a place on the LG Projects Wiki to further develop this work.

The comments the article received indicated that many of the pieces were in place. Most I knew about and this didn't surprise me. However, I want to build a solution rather than present a shopping list. That solution has to include various pieces of software, all playing together along with support. The pieces I see are:

- Audio conversion tools
- Audio editor
- Streamer
- Station automation
- Logging
- Transmitter

The first three items are pretty obvious and there are lots of choices out there. Logging is equally obvious. The other two items warrant further discussion once I get the basic concepts out.

## CUSTOMER BASE

You can look at this potential customer base in three

ways:

- Where you can get some money from
- What sort of fun code you can write
- What good political ends you can reach

The market is huge and very diverse. As a result of that diversity there are lots of choices of where you would like to fit into this project. For example, if you have a political ax to grind, helping your favorite political or religious cause get a radio voice or build a more efficient radio voice is likely your place. If, however, you just want to be in it for the money, there are tens of thousands of radio stations that could use your help and would pay for it.

Just looking at FM broadcast stations, there are 100 available frequencies on the band. In the US, these frequencies are allocated based on the signal strength of other stations on the same and adjacent frequencies. I don't have the numbers handy but that easily means thousands of stations in the US alone. Toss in AM broadcast and shortwave stations and you have the potential customer base for a product. If that isn't enough, Internet-only stations could use the same system.

Hopefully, if you are still reading, this is starting to make sense. I will describe the overall project idea and then talk about the specifics from above.

## AN OVERALL LOOK

Let me present how a traditional radio station works. By traditional, I mean how stations have worked for many years. That will make it easier to see where automation makes sense.

A typical station is composed of one or more studios and a transmitter facility. A studio is nothing more than a soundproof room with some audio equipment. The transmitter may be in the same building or at a remote location connected by a dedicated wire or radio link. While there is some room for talking about the transmitter link, I want to concentrate on the studio end.

Generally, one studio will be *live*. That means whatever is happening in that studio will be sent directly to the transmitter. The alternative to a live studio would be the station either re-broadcasting some external feed or something that they have pre-recorded. In any case, whenever the transmitter is on the air, there must be some source of program material. Additional studios are available to build the pre-recorded program material.

Each studio will likely contain one or more live microphones, multiple sources for pre-recorded material (CD players, turntable and tape decks) and something to record a program (tape recorder or mini-disk are the most common). Also included is a audio mixer board and monitor system so multiple sources can be mixed and edited.

Small stations will typically have one studio with a

person that queues and starts pre-recorded sources and makes live announcements that can include news and weather. For a typical music station, the majority of this person's time is spent waiting for the current song to end so they can queue the next one possibly inserting some commentary between songs. They may also be logging what they play so the station can pay the necessary royalties.

## STATION AUTOMATION

This is the most obvious piece of the system - replacing the live tedium of queuing CD tracks with an automated system. The most basic step would be to just save all the tracks on a disk in a computer system and allow the person to pre-select what they would like played during a specified time block. This is relatively trivial to do. We can, however, do a lot more than this.

Most stations develop a play list that the live DJs need to follow. This list includes songs that can be played along with guidelines for how often they can be played. Armed with that list, a program could easily make all the necessary decisions to offer what appears to be a random selection of music within the necessary constraints.

Next comes the announcements. They can be divided into:

- Commercials
- Information that will be supplied live or almost live
- Information that could be automated

Automation of commercials is not much different than the music play lists. The primary difference is that there will likely be specific times when a commercial must be broadcast. Nothing magic here--just another type of event to put into a scheduling program.

The live or almost live program material is that which must be put together by a person. For example, a news broadcast. This could either be done live by a person in the studio (or remotely over an Internet connection) or it could be pre-recorded. If the news was pre-recorded as a set of items then later news broadcasts could re-use the appropriate portions of a previous broadcast. In any case, there is still nothing particularly difficult about this--it is just another event to schedule.

I separated out information that could be automated because it is an additional project. That is, it does not have to be part of this original package. Two things come to mind here:

- Time announcements
- Weather

Both of these announcements are really nothing more than building a human voice announcement out of some digital data. Both have been done--it just becomes a matter of integration.

Some stations allow call-in requests. This means a human responds to a request, checks to see if the song

is available and hasn't been played more than is considered correct by the play list and then schedules it. This seems like a perfect place to use a web page. Listeners could place the requests and the software would appropriately adjust what was to be played. If the requested material was not available it could advise you of that fact or even order heavily requested material.

## TRANSMITTER

While a traditional transmitter is far away from the scope of this project, a transmitter option deserves mention. In my searching for FM broadcast transmitters for a project in Nicaragua I ran across a company that offers an FM broadcast transmitter on a PCI card. They have a new card on the way and I have offered to write the Linux driver for it.

If you saw this project as *just fun* or something that would work for your college dorm, a card such as this might be just the ticket. You could offer a web interface to select program material and broadcast to nearby FM receivers. Being a *radio guy* as well as a computer guy, this interests me a lot.

## OVERALL SCOPE

That covers all the pieces from the geek end. But, as I said in the beginning, I want to present a solution rather than a shopping list. That 100mw station in your basement that your family can listen to is not going to be a good commercial customer but tens of thousands of commercial stations certainly could be.

Integrating all the software necessary to offer a solution is the first part. Installation and support comes next. A radio station that has full-time employees running the station and advertisers paying hundreds or thousands of dollars for a commercial will quickly see the benefits of a system such as this. The biggest hurdle will be showing them that the solution will be supported. That is, that their station will not be off the air because they made this choice.

For a small station, this might mean knowing that they could call someone and have them come in and fix a problem. For a larger station it could mean training of on-site personnel. There are other levels of support including spare systems, shared servers and so forth. In other words, an assortment of different markets where the same software is offered but the support potential varies.

## WHAT NEXT?

That's up to you. I am excited about the project. Unfortunately, I have a "day job" which does not give me the necessary time to put all this together and even if it was together, I don't want to go into the software support business.

My hope is that the right people are out there that want to do the pieces. That is my real reason for writing this

article. I am happy to offer input, help set direction, offer a mailing list of discussion forum and even publicize the product. If you are interested in participating, write at [phil@ssc.com](mailto:phil@ssc.com) and let me know your interests. Or go out to the LG Projects Wiki and chime in. Maybe Linux-controlled radio will be here before you know it.

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

<http://www.linuxgazette.com/node/view/114>

## AUUG Corporate Members

as at 1st June 2003

- ◆ ac3
- ◆ Apple Computer Australia Pty Ltd
- ◆ Aptitude
- ◆ Australian Bureau of Statistics
- ◆ Australian Taxation Office
- ◆ British Aerospace Australia
- ◆ Bureau of Meteorology
- ◆ Cape Grim B.A.P.S
- ◆ Computer Associates
- ◆ Computing Services, Dept Premier & Cabinet
- ◆ Corinthian Industries (Holdings) Pty Ltd
- ◆ Cray Australia
- ◆ CSIRO Manufacturing Science and Technology
- ◆ CSIRO Telecommunications & Industrial Physics
- ◆ Curtin University of Technology
- ◆ Cybersource
- ◆ Deakin University
- ◆ Department of Land & Water Conservation
- ◆ Everything Linux & Linux Help
- ◆ Fulcrum Consulting Group
- ◆ IBM
- ◆ IP Australia
- ◆ KAZ Technology Services
- ◆ Land and Property Information, NSW
- ◆ LPINSW
- ◆ Macquarie University
- ◆ Multibase WebAustralis Pty Limited
- ◆ NSW Public Works & Services, Information Services
- ◆ Peter Harding & Associates Pty. Ltd.
- ◆ Powerhouse Museum
- ◆ Rinbina Pty. Ltd.
- ◆ Squiz Pty Ltd
- ◆ Sun Microsystems Australia
- ◆ Sydney Water Corporation
- ◆ Tellurian Pty. Ltd.
- ◆ The University of Western Australia
- ◆ Thiess Pty Ltd
- ◆ Uni of NSW - Computer Science & Engineering
- ◆ UNITAB Limited
- ◆ University of New England
- ◆ University of New South Wales
- ◆ University of Sydney
- ◆ University of Technology, Sydney
- ◆ Workcover Queensland



**CALL FOR PAPERS**  
**AUUG SECURITY SYMPOSIUM 2004**  
**20 FEBRUARY 2004,**  
**CANBERRA**

The AUUG Security Symposium provides a forum for discussion of security technologies, techniques and management.

Our society today is highly dependent on our almost pervasively interconnected systems. Hence we are also dependent upon the security of these systems. As Governments and private industry become increasingly aware of the vulnerability of our systems there is a growing requirement for security education and for practioners to share their knowledge for the greater good.

This symposium aims to fill a gap in the Australian conference scene between the high cost commercial conferences where attendees hear mainly marketing pitches and the academic-based research conferences. It is unashamedly for the practioner in the field who wants to share (or know about) how to secure their systems (be it a PC operating system, a huge network or a client server application).

The symposium will be a one day event with paper presentations. This will be the best opportunity in Australia in the coming year to meet, discuss and debate your ideas and experiences on information security.

The Call for Papers can be read in more detail at:  
<http://www.auug.org.au/events/2004/security/cfp.html>

The Security Symposium committee is:

Ben Elliston  
 Martin Schwenke  
 Stephen Rothwell

## Working With ACLs in FreeBSD 5.x

Author: Grzegorz Czapliński <[gregory@prioris.mini.pw.edu.pl](mailto:gregory@prioris.mini.pw.edu.pl)>

In this article I will describe how to create and configure unique access permissions on files and directories using Access Control Lists (ACLs).

ACLs provide an extended set of permissions for a file or directory. These permissions can be used in addition to the conventional UNIX permissions for files and directories. Standard UNIX file permissions provide read, write and execute access to three user classes:

- file owner
- file group
- others

ACLs are used to provide greater data access control for each file or directory. They enable you to define permissions for specific users and groups.

Every ACL has the following syntax:  
 [ACL tag]:[ACL qualifier]:[Access permissions]

ACL tag is a scope of the file permissions to the owner, group, others, specific users, specific groups or ACL's mask.

The ACL qualifier field describes the user or group associated with the ACL entry. It might be UID or user's name, GID or group's name, or empty.

Access permissions are the effective permissions for [ACL tag] and are specified as:

- r - read
- w - write
- x - execute

Entry types:

```
u::perm
permissions for the file owner
g::perm
permissions for the file group
o::perm
permissions for the others
u:UID:perm
permissions for the specific user identified by UID
u:username:perm
permissions for the specific user identified by
username
g:GID:perm
permissions for the specific group identified by
GID
g:groupname:perm
permissions for the specific group identified by
groupname
m::perm
maximum effective permissions allowed for specific
users or groups
```

The mask does not set the permissions for the file owner or others. It is used as a quick way to change effective permissions for all specific users or groups.

ACLs are part of UFS2 filesystem shipped with FreeBSD 5.0-RELEASE as an option or FreeBSD 5.1-RELEASE as the default filesystem during the installation. To check which filesystem you have on your system type:

```
# dumpfs /tmp | head -1
magic 19540119 (UFS2) time Fri Aug 15 19:23:30 2003
```

You must have ACL support compiled into the kernel too. Add:

```
options UFS_ACL          #Support for ACLs
```

to your kernel config compile and install a new kernel according to the instructions in the June 2003 Answerman column

(<http://ezine.daemonews.org/200306/answerman.html#kbuild>)

To enable ACLs on a partition, after newfs(1)'ing it issue

the commands:

```
# tune2fs -a enable /dev/dals1e
# mount /dev/dals1e /mountpoint
# mount | grep acl
/dev/dals1e on /mountpoint (ufs, local, soft-
updates, acls)
```

This indicates that soft-updates and acls are enabled on the /dev/dals1e partition mounted under /mountpoint.

The other way to check if ACLs are enabled is to use tune2fs(1) command:

```
# tune2fs -p /dev/dals1e
tune2fs: ACLs: (-a)          enabled
tune2fs: MAC multilabel: (-l) disabled
tune2fs: soft updates: (-n)  enabled
tune2fs: maximum blocks per file
in a cylinder group: (-e)   2048
tune2fs: average file size: (-f) 16384
tune2fs: average number of files
in a directory: (-s)       64
tune2fs: minimum percentage
of free space: (-m)        8%
tune2fs: optimization preference: (-o) time
tune2fs: volume label: (-L)
```

Before I show some examples please read the manpages for setfacl(1) and getfacl(1). The commands and their output below are separated by one empty line for clarity. On my test system I have a user called acl and he belongs to wheel group. When you see touch(1) command in an example, that means I recreated a file after it was removed.

#### Modifying ACLs

Create an empty file:

```
% umask 027
% touch file.txt
% ls -l file.txt
-rw-r----- 1 acl wheel 0 Aug 5 22:35 file.txt
```

```
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
group::r--
other::---
```

The file.txt is a normal file without any ACL permissions set yet.

#### ADD AN ACL PERMISSIONS:

```
% setfacl -m u::rw,g::r,u:gregory:rw file.txt
(! Change "gregory" to some other username.)
% ls -l file.txt
-rw-rw-----+ 1 acl wheel 0 Aug 5 22:41 file.txt
```

The little "+" at the end of access rights column indicates that the file has ACL set.

```
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
user:gregory:rw-
```

```
group::r--
mask::rw-
other::---
```

This command shows that owner has read/write access, group has read access, and user gregory has read/write access. I have to point now that the mask indicates the maximum permissions for user gregory. If the command was (set the mask - "m::r"):

```
% setfacl -m u::rw,g::r,u:gregory:rw,m::r file.txt
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
user:gregory:rw-          # effective: r--
group::r--
mask::r--
other::---
```

user gregory would have read/write access, but the mask would downgrade the effective access rights to read only.

There is an "-M" switch that is used to set and modify the ACL entries. The information about actual ACLs are kept in a file (in this example acls.txt).

```
% touch file.txt
Create acls.txt file which looks like:
% cat acls.txt
u:bin:rwX
```

```
% setfacl -M acls.txt file.txt
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
user:bin:rwX
group::r--
mask::rwX
other::---
```

In the last example ACL entry for user bin was specified in a file acls.txt.

#### RECALCULATING AN ACL MASK

The ACLs look as above (the last getfacl(1) command), issue a command:

```
% setfacl -m u::rw,g::r,u:bin:rw file.txt
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
user:bin:rw-
user:gregory:rw-
group::r--
mask::rw-
other::---
```

Now, users gregory and bin have read/write access, and the mask has been "group" ACL entries in the resulting ACL.

If the last command was:

```
% setfacl -n -m u::rw,g::r,u:bin:rw file.txt
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
user:bin:rw-          # effective: r--
user:gregory:rw-     # effective: r--
group::r--
mask::r--
other::---
```

the mask would not get recalculated (switch -n). Effective rights for users gregory and bin would be read only.

## DELETING AN ACL

To delete an ACL entry for user bin do:

```
% setfacl -n -x u:bin:rw file.txt
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
user:gregory:rw-     # effective: r--
group::r--
mask::r--
other::---
```

The entry for user bin was deleted. If you want the mask not to get recalculated, remember to use the "-n" switch. If you didn't use it, the mask would be read/write now, effectively changing permissions for user gregory to read/write.

To remove permanently ACL from a file issue:

```
% setfacl -bn file.txt
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
group::r--
other::---
% ls -l file.txt
-rw-r----- 1 acl wheel 0 Aug 5 23:08 file.txt
```

Compare the above with that:

```
% setfacl -b file.txt
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
group::r--
mask::r--
other::---
```

In the next example, setfacl(1) command is able to change permissions for all user classes - owner, group, others.

```
% umask 027
% touch file.txt
% ls -l file.txt
```

```
-rw-r----- 1 acl wheel 0 Aug 5 23:13 file.txt
```

```
% setfacl -m u::rw,g::r,o::r,u:gregory:rw file.txt
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
user:gregory:rw-
group::r--
mask::rw-
other::r--
alphax% ls -l file.txt
-rw-rw-r--+ 1 acl wheel 0 Aug 5 23:12 file.txt
More interesting example:
% touch file.txt
% ls -l
total 0
-rw-r----- 1 acl wheel 0 Aug 5 23:24 file.txt
```

```
% chmod 660 file.txt
% ls -l
total 0
-rw-rw---- 1 acl wheel 0 Aug 5 23:24 file.txt
```

```
% setfacl -m u::rw,g::r,o::r,u:gregory:rw file.txt
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
user:gregory:rw-
group::r--
mask::rw-
other::r--
```

```
% ls -l
total 2
-rw-rw-r--+ 1 acl wheel 0 Aug 5 23:25 file.txt
```

```
% chmod 644 file.txt
% ls -l
total 2
-rw-r--r--+ 1 acl wheel 0 Aug 5 23:25 file.txt
```

```
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
user:gregory:rw-     # effective: r--
group::r--
mask::r--
other::r--
```

The last setfacl(1) command set the access rights as follows:

```
user::rw-
user:gregory:rw-
group::r--
mask::rw-
other::r--
```

Then I changed explicitly access rights with chmod(1) command:

```
% chmod 644 file.txt
```

and the access rights reappeared as:

```
user::rw-
user:gregory:rw-     # effective: r--
group::r--
```



```
mask:r--
other:r--
```

Note, the mask is closely associated with group access rights. Changing Unix access rights with `chmod(1)`, you also change the mask value.

Consider this scenario:

```
% touch file.txt
% setfacl -m u::rw,g::rw,o::r,u:gregory:rw file.txt
% ls -l file.txt
-rw-rw-r--+ 1 acl wheel 0 Aug 6 20:19 file.txt
```

```
% setfacl -m m:r file.txt
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
user:gregory:rw-      # effective: r--
group::rw-           # effective: r--
mask:r--
other:r--
```

```
% ls -l file.txt
-rw-r--r--+ 1 acl wheel 0 Aug 6 20:20 file.txt
```

Changing the mask value, does change group access rights.

If you see a file with a magic "+" at the end of access rights column, check it with `getfacl(1)`.

## COPYING ACL ENTRIES

```
% touch file.txt
% setfacl -m u::rw,g::r,u:gregory:rw file.txt
% getfacl file.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
user:gregory:rw-
group::rw-
mask::rw-
other:r--
% getfacl file.txt | setfacl -b -n -M - file1.txt
```

```
% getfacl file.txt file1.txt
#file:file.txt
#owner:1009
#group:0
user::rw-
user:gregory:rw-
group::rw-
mask::rw-
other:r--
```

```
#file:file1.txt
#owner:1009
#group:0
user::rw-
user:gregory:rw-
group::rw-
mask::rw-
other:r--
```

## CREATING DEFAULT ACLs

Default ACL entries provide a way to propagate ACL information automatically to files and directories. New

files and directories inherit ACL information from their parent directory if that parent has an ACL that contains default entries. You can set default ACL entries only on directories.

Example:

```
% umask 027
% mkdir dir
% ls -l
total 2
drwxr-x--- 2 acl wheel 512 Aug 6 11:50 dir
```

```
% getfacl dir
#file:dir
#owner:1009
#group:0
user::rwx
group::r-x
other:----
```

Before you set any default ACL entries for users or groups, you must set default ACL entries for owner, group, other, and ACL mask.

Consider this:

```
% setfacl -m u::rwx,m::rwx,g::rx,o::rx dir
% getfacl dir
#file:dir
#owner:1009
#group:0
user::rwx
group::r-x
mask::rwx
other:r-x
% setfacl -dm u:gregory:rwx,m::rwx dir
```

```
setfacl: acl_set_file() failed for dir: Invalid argument
The correct order is:
```

```
% setfacl -dm u::rwx,m::rwx,g::rx,o::rx dir
```

1. Set default ALC entries for directory owner, group, others and the mask.

```
% getfacl -d dir
#file:dir
#owner:1009
#group:0
user::rwx
group::r-x
mask::rwx
other:r-x
```

To view default ACLs issue `getfacl(1)` with the

"-d" switch.

```
% setfacl -dm u:gregory:rwx,m::rwx dir
```

2. Set default ALC entry for user gregory. To see the effect of default ACLs on subdirectories issue the following commands:

```
% mkdir dir/subdir
% getfacl -d dir
#file:dir
#owner:1009
#group:0
user::rwx
user:gregory:rwx
group:r-x
```

```
mask::rwx
other::r-x

% getfacl -d dir/subdir
#file:dir/subdir
#owner:1009
#group:0
user::rwx
user:gregory:rwx
group::r-x
mask::rwx
other::r-x
```

The subdir directory successfully inherited default ACL entries from its parent.

Suppose, you want to set default ACL entries for additional user bin:

```
% setfacl -dm u:bin:rwx,m::rwx dir
% getfacl -d dir
#file:dir
#owner:1009
#group:0
user::rwx
user:bin:rwx
user:gregory:rwx
group::r-x
mask::rwx
other::r-x
```

```
% getfacl -d dir/subdir
#file:dir/subdir
#owner:1009
#group:0
user::rwx
user:gregory:rwx
group::r-x
mask::rwx
other::r-x
```

That new default ACL entries for additional user bin are not visible on dir/subdir as the directory was created before the ACL entry for bin was set.

To see the effect of default ACLs on files, create a file beneath the dir directory:

```
% touch dir/file.txt
% ls -l dir/file.txt
-rw-r-----+1 acl wheel 0 Aug 6 12:14
dir/file.txt
```

```
% getfacl dir/file.txt
#file:dir/file.txt
#owner:1009
#group:0
user::rw-
user:bin:rwx # effective: r--
user:gregory:rwx # effective: r--
group::r-x # effective: r--
mask::r--
other::---
```

The setfacl(1) manual states: "Currently only directories may have default ACL's."

## DELETING DEFAULT ACLS

To delete default ACLs on directories, use setfacl(1) with the "-k" switch:

```
% setfacl -k dir
```

```
% getfacl -d dir
#file:dir
#owner:1009
#group:0
```

```
% getfacl dir
#file:dir
#owner:1009
#group:0
user::rwx
group::r-x
other:---
```

To delete a default ACL entry for user bin do:

```
% mkdir dir
% setfacl -dm u::rwx,m::rwx,g:rx,o:rx dir
% setfacl -dm u:gregory:rwx,u:bin:rwx,m::rwx dir
```

```
% getfacl -d dir
#file:dir
#owner:1009
#group:0
user::rwx
user:bin:rwx
user:gregory:rwx
group::r-x
mask::rwx
other::r-x
```

Create acls.txt file which looks like:

```
% cat acls.txt
u:bin:rwx

% setfacl -dX acls.txt dir
alpha% getfacl -d dir
#file:dir
#owner:1009
#group:0
user::rwx
user:gregory:rwx
group::r-x
mask::rwx
other::r-x
or simply type:
% setfacl -d -x u:bin:rwx dir
```

## THINGS TO REMEMBER

setfacl(1) always recalculates the ACL mask to allow maximum effective permissions for every ACL entry, unless the "-n" switch is used.

If you use the chmod(1) command to change the file group owner permissions on a file with ACL entries, both the file group owner permissions and the ACL mask are changed to the new permissions. Be aware that the new ACL mask permissions may change the effective permissions for additional users and groups who have ACL entries on the file.

ACLs are developed on FreeBSD-5.X line. Things may change rapidly from time to time. If you use any of FreeBSD-5.X follow freebsd-current mailing list. On the August 3rd, Robert Watson sent a message saying that he has modified the semantics of the application of the default ACL in combination with the umask.

Because of the development process you may notice

differences in output, therefore I hope my article will give you good insight on ACLs.

ACLs are a good way to improve filesystem security on your server. They are flexible and as a new option shipped with the UFS2 filesystem make FreeBSD a more powerful system than ever.

Special thanks go to Jason McIntyre, Joe Warner and other proofreaders.

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

*<http://ezine.daemonnews.org/200310/acl.html>*

---

# Mandrake 9.1 - Raising the Bar for GNU/Linux Distribution Usability

Author: Andrew D. Balsara <[andrebal@a mailingaddress.org](mailto:andrebal@a mailingaddress.org)>

## REVIEW



Mandrake 9.1 is ready, and I have also received an email from [MandrakeClub](#) indicating that 9.1 boxed CD sets are now available for pre-order (with a hefty discount for club members). Sounds like a good time to put myself in the seat of a first-time Linux user and dissect this long-awaited release. I had rehearsed the installation of Mandrake Linux 9.1 final using 9.1 RC2, and this time I decided to explore in detail the hardware and software configuration that I would consider adequate for a comfortable Mandrake desktop experience. Kind of, a Mandrake 9.1 Certify-It-Yourself PC if you see what I mean...

See also the previous installments: [Part 1](#) [Part 2](#) [Part 3](#) [Part 4](#) [Part 5 \(PPC\)](#) [Part 6](#)

## PROGRAMMES AND SPONSORSHIP

[Timesavers](#) Thank you for supporting [DistroWatch Advertise](#)

## OPINIONS AND FACTS

It's inevitable that when writing an article some opinions creep into the text, even though I try to write about observations and facts while avoiding impressions and opinions. For example, on the matter of desktop GUIs (Graphical User Interfaces), while I do like to use both KDE and GNOME applications, I rather tend to use KDE + Konqueror + KMail + KOffice for my daily work.

This is a matter of taste and habit, it is not a choice based on a detailed comparison between KDE and GNOME and their equivalent applications (BTW this is another reason why I avoid comparing GUIs, Operating Systems, distributions, etc: such comparisons can easily become too subjective to have any value). So, what works for me might not work for a large number of Linux users, and vice-versa. I hope the pictures speak for themselves, though.

## Screenshots 2 & 3: A KDE application - Kget

### INSTALLING MANDRAKE LINUX 9.1

I previously reviewed the various Mandrake 9.1 Betas and also RC1, and between each release I could observe the progress being made with DrakX, the Mandrake Linux installer program. I am happy to say that DrakX got quite a bit of polish for 9.1 final: not only does it have an extremely professional look, but it worked flawlessly, with the single exception of configuring my dual-monitor setup (DrakX configured X Window for a single monitor and then I updated the /etc/X11/XF86Config-4 configuration file manually after the first boot).

### Screenshot 4 & 5: Mandrake 9.1 installation - partitioning and package selection

I did not time the installation but it must have taken about 50 minutes total, including two small pauses for drinking a cup of tea, and watching the news. Most of the time during installation was spent by DrakX automatically installing the selected packages, and this ran unattended. This time I had reserved 10GB of free disk space for my Mandrake Linux 9.1 installation, which I partitioned as follows:

Partition	Mount point	Size	Filesystem
hda1	/boot	32MB	ext2
hda5	/	4GB	reiserFS
hda6	swap	250MB	
hda7	/tmp	500MB	reiserFS
hda8	/var	500MB	reiserFS
hda9	/home	4.6GB	reiserFS

Table 1: Partitioning example for Mandrake Linux 9.1

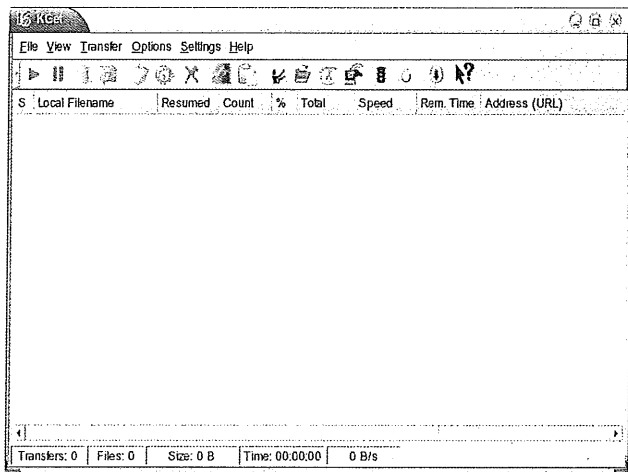
For all my previous reviews I had used a single 5GB partition, but this is more of a final "production" installation and I feel more comfortable with this partitioning scheme.

### Screenshot 6: Mandrake 9.1 installation - hardware configuration

After installing all the packages DrakX still needs a few simple steps to configure the system before rebooting. The many choices may seem daunting at first but there is no reason to be scared. Each step is accompanied by extensive Help instructions, and the defaults are in most cases adequate and can be changed later.

### INITIAL CONFIGURATION

I prefer to boot directly into the Display Manager and Mandrake 9.1 has a great choice in this respect: one can choose from KDM (KDE style), GDM (GNOME style), MdkKDM (Mandrake style) or XDM (plain X Window style). By default Mandrake Linux 9.1 sets up MdkKDM, but I prefer KDM, the original (and more



sophisticated) KDE Display Manager. That's easy to configure with some excellent utilities included in Mandrake Linux 9.1:

1. From the console command line or a terminal, use `urpmi` to install the KDM package:  
`urpmi kdebase-kdm`
2. Select your favorite Display Manager with `drakdm`  
`drakdm`
3. On your next boot, you will be presented with KDM's fully customizable login.
4. Open the KDE Control Center, find the *Configuration / KDE / System / Login Manager* option and configure KDM to your taste.

#### Screenshot 7 & 8: Mandrake 9.1 - login and menudrake configuration

Another small issue I have with the default setup is that the KDE Calculator was not included in the *Office / Accessories* menu. I found KCalc under the *Applications / Sciences / Mathematics* menu. Again, this is easily changed using the Menudrake tool (found under *Configuration / Other*). I already discussed in my previous articles some of the more common Linux desktop customizations, so here is a typical screenshot, showing how I prefer to work:

#### Screenshot 9: Mandrake 9.1 - a desktop example

And here is a short description of some visual elements that have been customized:

- The main customization which is usually performed at installation is the choice of XFree86 driver, resolution and color depth. As obvious from the picture above, I am using the Xinerama extension which allows the use of two or more monitors. Users with a single monitor can rely on DrakX and not worry.
- The second choice is the choice of GUI. Both GNOME and KDE offer excellent usability, but I find KDE's operation simpler to grasp for beginners. Mandrake Linux 9.1 includes a few other window managers but they are more suited to specific user tastes or needs.

- KDE 3.1's Keramik theme with a slightly modified Keramik Emerald color scheme (I call it Galaxy Bronze, because it also works well with Galaxy).
- Verdana font and variations, with anti-aliasing turned on.
- A background pixmap with a slight gradient for the KDE Panel.

#### Screenshot 10: Mandrake 9.1 - KDE's kicker

- Four customized Virtual Desktops with distinct names.
- Custom background color (pale green) for the Digital Clock.
- Personalized simple background, tiled.

#### NON-VISUAL ELEMENTS

I usually also personalize some non-visual elements for my desktop:

- Sounds can be configured using the *Configuration / KDE / LookNFeel / System Notifications* control panel. Note that each KDE application can have its own set of sounds.
- Mouse and keyboard can be configured using the respective controls panels in *Configuration / KDE / Peripherals*.

#### Screenshot 11 & 12: Sounds and mouse configuration panel; note the various tabs allowing total customization of the mouse device

If this smorgasbord of customization options seems daunting at first, I consider it a plus, specially since Mandrake Linux 9.1 provides extremely good defaults to start with. The Galaxy theme and default font rendering setup provide a pleasing desktop out-of-the-box.

Did I mention that networking was configured automatically by DrakX (I use DHCP)? It seems that I am beginning to take things for granted with this new release, but really I didn't even worry about checking the settings for my network configuration. Network configuration is done (as root) through the Mandrake Control Center. Similarly, my network printer was perfectly detected and configured without any input on my part. I just had to launch the printer installation option in the Mandrake Control Center:

#### Screenshot 13 & 14: Network and printer configuration panel


#### THE MANDRAKE PC

Hardware manufacturers have long been standardizing on features and ergonomic details for their products (in part because of the dominance of a single commercial desktop OS and a single microprocessor company). GNU/Linux has benefited from this not only because it

is easier to write a single driver for a specific hardware than many different ones, but also because having clear specifications laid down avoids having to reverse engineer proprietary drivers.

**Picture 1: These connectors have become a standard on PCs**

The latest standard basic PC specifications consist of the PC 2001 System Design Guide (check [www.pcdesguide.org](http://www.pcdesguide.org) ) and many related documents. One can also add the USB 1.1 and 2.0 specifications, the ATA specifications, the IEEE1394 specifications and many others. All these are published standards which are freely available to the public and provide a wealth of information to Linux developers.



**Picture 2: The iMac**

Presently I would say that any 800MHz or faster x86 CPU is adequate for a Mandrake Linux 9.1 personal workstation. So, even an inexpensive AMD Athlon XP 1700 (which runs @ 1.47GHz, by the way) or Intel Celeron 1.7GHz is overkill. I have not tested the latest VIA CPUs but I have some reason to believe that their latest C3 Nehemiah running @ 1GHz would also provide satisfactory performance, although it is no match for the other two CPUs for heavy processing. Historical detail: Linus Torvalds developed the first kernel on a 386DX 33MHz machine circa 1991. CPUs and Linux have come a long way since...

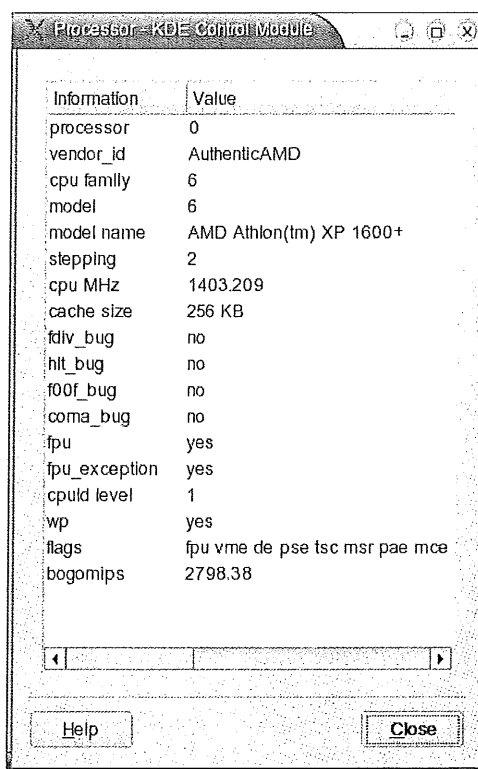
### SYSTEM DESIGN GUIDE

It can be said that the role of a Linux distribution vendor is to integrate all these different layers, but the stumbling block in this process is hardware compatibility. From that point of view Mandrake Linux 9.1 does quite well: it manages to handle the widest range of hardware of any Linux distribution I have yet to try.

So let's take a look at what makes a Mandrake Linux 9.1 PC tick and whether there are any issues with currently available hardware.

### CHOICES AND RECOMMENDATIONS

I am a Mac/PC agnostic. I have used Macs since the very first Mac 128K (anybody remember the WriteNow word processor?), and more recently I reviewed Mandrake PPC 9.1 Beta (the final release should be coming out Real Soon Now) on an iMac 350MHz with surprisingly good results. Whatever my feelings, the reality is that Macs now occupy a niche market representing perhaps 2-3% of the total microcomputer market and this share is shrinking. I had high hopes that Digital/Compaq/HP would turn the Alpha into a mass-market product but that never happened and seems unlikely now. It seems we are stuck with x86 CPUs for some more time! Right now I feel engineers have done a good job performance-wise, so that the choice of CPU architecture and the old debates about RISC vs. CISC are becoming almost irrelevant\*.

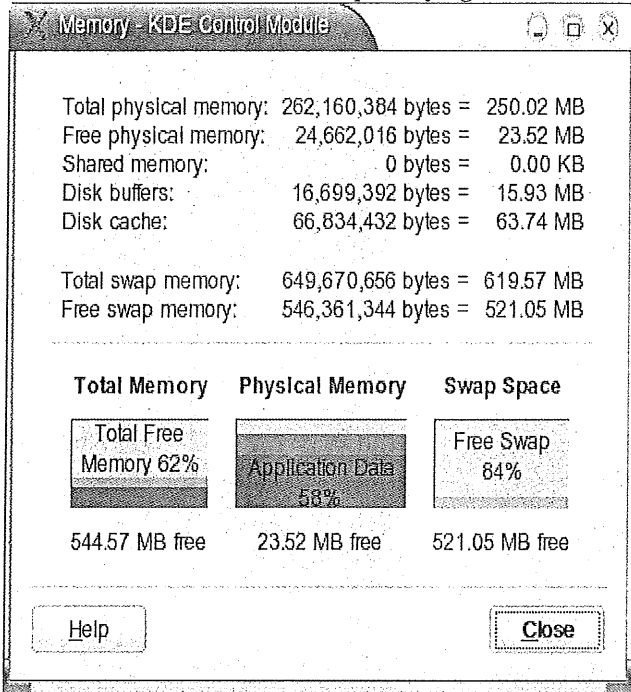


**Screenshot 15: KDE control module displaying CPU**

**information**

The mainboard is the next component that determines the performance and reliability of a computer, but here too enormous progress has been done in terms of integrating features and quality and lowering costs. Most OEMs or do-it-yourself mainboards are fully compatible with the latest Linux kernel included with Mandrake Linux 9.1.

Another common issue is that of RAM; given the low prices of RAM these days I would recommend any Linux Mandrake 9.1 user to add or upgrade to a 256MB stick of RAM in his machine and stop worrying.



**Screenshot 16: KDE control module displaying memory information**

Then there is the issue of video card choice. Unfortunately the XFree86 project tends to lag behind significantly in terms of support for the latest video chipsets and 3D features. My recommendation is to use a video card one or two generations older than the state-of-the-art. These are inexpensive and usually will still provide plenty of service under Linux. Practically all the AGP cards commonly available on the market are supported by XFree86 4.3. NVidia cards have the additional advantage of being supported by nVidia drivers, developed independently of the XFree86 group. Also, ATI is offering its own Linux drivers for the Radeon 8500 and above cards. Integrated video is usually not a good choice, due to poor performance and precarious driver support (with the exception of the latest nVidia chipsets with integrated GeForce 4 MX video).

Similarly, when it comes to notebooks the latest Linux kernel seems to have incorporated most of the advances in chipsets and power management, and due to thermal dissipation requirements notebooks tend to use slightly

simpler video hardware, which is usually well-supported under XFree86. Regarding disk space, as mentioned above I would consider 10GB adequate for Mandrake Linux 9.1, with some room left for future expansion. Performance is not issue with most modern IDE hard disks.

An accessory that has become indispensable is the CD-RW drive, which nowadays costs little more than its read-only CD-ROM counterparts. Some models are coupled with a DVD drive and more recently with DVD-R and DVD-RW capabilities, but these are not mainstream yet.

The rest of the system is non-critical but I would note the following:

- A wheel mouse is a good choice for general Linux desktop use, and wheel mice are cheap.
- Some sound cards are difficult to configure or not fully supported, but even an inexpensive CMI 8738 PCI sound card provides good sound support under Linux. Mandrake Linux 9.1 uses the ALSA drivers and as far as I could test these drivers support the vast majority of sound cards and chips on the market. And the best part is that DrakX is able to identify the sound chips and configure modules and sound settings automatically.
- A good monitor is a must for extended usage periods. And a resolution of 1024x768, the standard for the last few years, falls slightly short of the ideal resolution for Linux Mandrake 9.1 desktop use. The next step is 1280x1024, available with most 17" monitors and LCD screens. Notebooks are now available with resolutions of 1400x1050 and above, a big plus for Linux users.

Below is a table describing hardware that I would consider adequate for Mandrake Linux 9.1 desktop use. By no means is this a selection of the "best" possible hardware for GNU/Linux, but rather some recommendations on inexpensive hardware that is currently available and easy to purchase.

Hardware component	Adequate performance level	Available off the shelf
CPU	> 800MHz x86	Athlon XP 1700+ or Celeron 1.8GHz
Mainboard	Any	Any with the following chipsets: - SiS 745 - SiS 648 - SiS 746 - SiS 65x - KT266A - i845 (all) - KT333 - KT400 - nVidia nForce2
RAM	> 256 MB	2 x 128MB or 1 x 256MB

Hardware component	Adequate performance level	Available off the shelf
Video	Any TnT2 or Radeon card	Any of the following types: - Radeon 7000 - nVidia (all) - Radeon 7500 - Radeon 8500 - Radeon 9000
Hard Disk	> 10GB free space	Any > 20GB, Maxtor, Hitachi, WD
Optical drive	> 12x CD-RW or "combo" drive	48x CD-RWs are commonplace
Monitor	1280x1024 or better @ 75Hz	17" monitor or LCD
Network	Any Fast Ethernet (100BaseTx)	Commodity item
Other	Wheel mouse, PCI sound card, etc	Commodity items

Table 2: An adequate configuration for Mandrake Linux 9.1 is easily put together

The above components put together, either in a homemade computer or in a commercial product, would constitute a reasonably fast and very usable Mandrake Linux 9.1 Workstation. The only thing missing would be a "Linux Certified" (or "Mandrake Linux 9.1 Inside") label to glue on the front of the case! Similarly, a good many modern notebooks are perfectly usable with Mandrake Linux 9.1.

#### PERIPHERALS

It's perhaps in the lack of drivers for peripherals that GNU/Linux is somewhat lacking, because there is simply no way that Linux developers can track in real time the release of new products in the market. There are just too many one-of-a-kind peripherals and unfortunately hardware vendors are still concentrating all their efforts on the dominant commercial OS. This is slowly changing, though. Note that there is only so much that MandrakeSoft can do in this case; a distribution vendor only integrates drivers from the many Linux developers around the world, but cannot be expected to develop them for every single piece of hardware available.

My advice here would be to take it case by case. Here are some more common peripherals and the kind of support that can be expected in this new Mandrake Linux 9.1 release:

Peripheral	Linux support	Mandrake Linux 9.1
Digital cameras	Good	Open GTKam; check the Camera /Add camera menu option
Scanners	Good	Depends on <u>SANE</u> (Scanner Access Now Easy)
USB peripherals	Good	Most common peripherals supported; check <u>Linux USB overview</u>
Printers	Average	Support is OK for B/W, but color is a notch below commercial drivers
Wi-Fi, Gigabit	Good	Known brands well supported
Graphics tablets	Poor	Few applications, lack of drivers
Webcams	Average	Some work, some don't; low image quality
TV Cards	Excellent	DrakX recognizes and configures many TV cards
One-of-a-kind gadgets	?	Depends...

Table 3: Peripheral support under Mandrake Linux 9.1

One last note is that for GNU/Linux users, it is safer to buy a mainstream peripheral than a state-of-the-art one. Mainstream hardware is usually better tested and most kinks have been worked out... and one can check beforehand if a Linux driver is available for it!

#### MANDRAKE LINUX 9.1: NOT JUST A FEW CDS

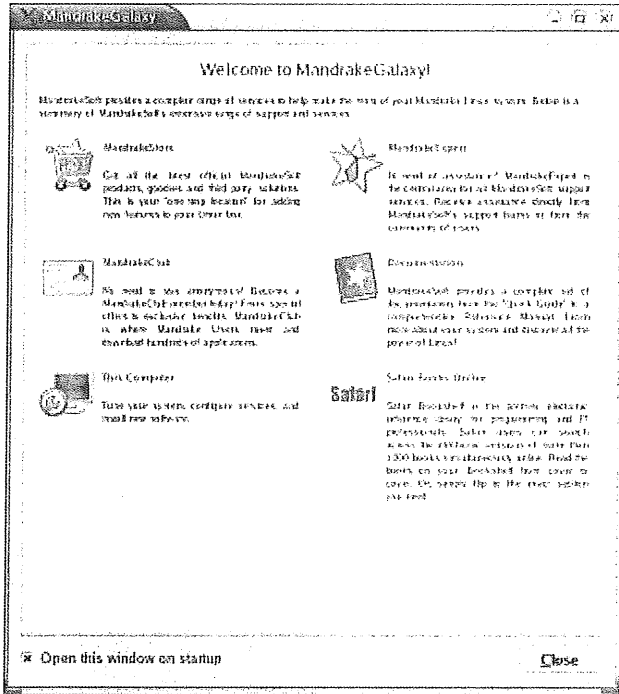
One of the things that makes Mandrake Linux 9.1 a great distribution is not just the fact that you can find so much quality Free software crammed in three CDs: it's also the information, support and community that has formed around it, as well as the additional services provided by MandrakeSoft over the Web. This may sound like advertising\*\*, but I believe the best way to have access to the various MandrakeSoft services and a good channel for communication is to become a MandrakeClub member.

MandrakeClub is a membership-based service from MandrakeSoft that is entirely available over the Internet. Not only do they provide a good level of support, but MandrakeClub members have access to even more pre-compiled packages for Mandrake Linux 9.1 - for free! Let's call it the topping on the cake.

#### CONCLUSION\*\*\*

Is Mandrake Linux 9.1 usable, even for first time Linux beginners? Most certainly! I could not find any serious flaw or obstacle either in the installation program, the default desktop configuration, the level of hardware support, or in the various applications included in the GNOME or KDE menus. On the contrary, the Galaxy theme presents an aesthetically pleasing desktop, and the variety and usability of the included applications is impressive, without being overwhelming. And all this works quite well on "average" hardware, without requiring any special investment.





**Screenshot 17: The Mandrake Galaxy tips and links helps beginners get started**

There is a growing feeling among I.T. professionals and particularly among Linux people that Linux is finally coming of age. I certainly feel that way when it comes to Mandrake Linux 9.1. This is the first of a wave of GNU/Linux distributions that have the potential to gain a considerable market and mind share among new users with no previous experience in Linux. And for those that have already been using Linux for some time, it probably means that the disk partition(s) holding other commercial operating systems can be erased and put to good use, at last!

**Note:** I recently browsed through a rather nice website for Linux beginners: [www.tinyminds.org](http://www.tinyminds.org). They have good practical tips for Linux users in their Cheat Sheets section. Certainly worth a check.

\* But the forthcoming launch of a 64-bit architecture on top of x86, as AMD has recently demonstrated (the processor will initially be called Opteron), could lead to exciting developments in the GNU/Linux world. MandrakeSoft is ready, having announced the immediate availability of [Mandrake Linux 9.0 for x86/64](#). I suppose that before the Opteron launch in April, MandrakeSoft will have ported its Mandrake Linux 9.1 release to this new architecture.

\*\* I am not an employee of [MandrakeSoft](#) or of [Distrowatch](#) and my reviews are written completely independently, on my own not-so-free time.

\*\*\* This article also concludes my series on Mandrake Linux 9.1; my next distribution review is in the works, so stay tuned @ [Distrowatch](#) !  
Copyright Notice

Copyright (C) 2003 Andrew D. Balsa  
Verbatim copying and distribution of this article is permitted in any medium, provided this copyright notice is preserved.

Review Notes	
Date	30 March 2003
Author	<a href="#">Andrew D. Balsa</a>
About the Author	Andrew D. Balsa is a Linux software developer and I.T. consultant based in Hong Kong. He is also the author of the Linux Benchmarking HOWTO.

Test Configuration (Desktop)	
Processor	Celeron 800 MHz
Motherboard	Abit BE-6 BX mainboard
Graphics Card	nVidia TNT64 Video card with 16MB SDRAM
Memory	256MB PC-100 SDRAM
Hard Disk	IBM 10GB DTTA-371010 (5GB used by NTFS filesystem)
Other Storage	5x DVD-ROM drive Panasonic SR-8583A
Network Card	Digital 21143
Monitor	KDS 21" CRT monitor 1280 x 1024 @ 75Hz
Sound Card	Generic CMedia 8738 PCI
Other	33.6K ISA Fax Modem (hardware), BT848 PCI TV Card, Acer Prisa 310U USB scanner, generic USB graphics tablet, Microsoft PS/2 IntelliMouse Web

Mandrake Specifications	
Version	9.1
Release Date	25 February 2003
Kernel	2.4.21-pre
GUI	XFree86 4.3.0
File Systems	ReiserFS ext3 ext2 JFS FAT ISO9660 XFS
Minimum Requirements	<ul style="list-style-type: none"> <li>i586 processor</li> <li>64MB RAM recommended, 32MB RAM for text install</li> <li>800MB recommended, minimum 500MB hard disk space</li> <li>CD-ROM or floppy drive</li> </ul>
Platforms	Pentium and compatible processors, AMD processors

Features	<ul style="list-style-type: none"> <li>• Mandrake Control Center</li> <li>• NTFS partition resizer</li> <li>• 100% Free Software</li> </ul>
Bundled Software	<ul style="list-style-type: none"> <li>• Apache 2.0.44</li> <li>• GCC 3.2.2</li> <li>• Gimp 1.2.3</li> <li>• glibc 2.3.1</li> <li>• GNOME 2.2</li> <li>• GTK+ 2.2.2</li> <li>• KDE 3.1</li> <li>• Mozilla 1.3</li> <li>• Perl 5.8.0</li> <li>• Python 2.2.2</li> <li>• Samba 2.2.7a</li> <li>• XFree86 4.3.0</li> <li>• Xmms 1.2.7</li> </ul>
Price	<ul style="list-style-type: none"> <li>• Standard Edition (US\$39.00)</li> <li>• PowerPack Edition (US\$69.00)</li> <li>• ProSuite Edition (US\$199.00)</li> </ul> <p>Available from <b>Mandrake Store</b></p>

# CyberInsecurity: The Cost of Monopoly

## (Part 1 of 2)

Authors: Dan Geer, Rebecca Bace, Peter Gutmann, Perry Metzger, Charles P. Pfleeger, John S. Quarterman, Bruce Schnerier

### INTRODUCTION BY COMPUTER & COMMUNICATIONS INDUSTRY ASSOCIATION

No software is perfect. This much is known from academia and every-day experience. Yet our industry knows how to design and deploy software so as to minimize security risks. However, when other goals are deemed more important than security, the consequences can be dangerous for software users and society at large.

Microsoft's efforts to design its software in evermore complex ways so as to illegally shut out efforts by others to interoperate or compete with their products has succeeded. The monopoly product we all now rely on is thus both used by nearly everyone and riddled with flaws. A special burden rests upon Microsoft because of this ubiquity of its product, and we all need to be aware of the dangers that result from reliance upon such a widely used and essential product.

CCIA warned of the security dangers posed by software monopolies during the US antitrust proceeding against Microsoft in the mid and late 1990's. We later urged the European Union to take measures to avoid a software monoculture that each day becomes more susceptible to computer viruses, Trojan Horses and other digital pathogens.

Our conclusions have now been confirmed and amplified by the appearance of this important report by leading authorities in the field of cybersecurity: Dan Geer, Rebecca Bace, Peter Gutmann, Perry Metzger, John S. Quarterman, Charles Pfleeger, and Bruce Schneier.

CCIA and the report's authors have arrived at their conclusions independently. Indeed, the views of the authors are their views and theirs alone. However, the growing consensus within the computer security community and industry at large is striking, and had become obvious: The presence of this single, dominant operating system in the hands of nearly all end users is inherently dangerous. The increased migration of that same operating system into the server world increases the danger even more. CCIA is pleased to have served as a catalyst and a publisher of the ideas of these distinguished authorities.

Over the years, Microsoft has deliberately added more and more features into its operating system in such a way that no end user could easily remove them. Yet, in so doing, the world's PC operating system monopoly has created unacceptable levels of complexity to its software, in direct contradiction of the most basic

Resources	
Home Page	<a href="http://www.mandrakelinux.com">www.mandrakelinux.com</a>
Quick Links	<a href="#">MandrakeSoft</a> <a href="#">MandrakeStore</a> <a href="#">MandrakeClub</a> <a href="#">MandrakeExpert</a> <a href="#">MandrakeBizCases</a> <a href="#">MandrakeForum</a> <a href="#">MandrakeUser</a> <a href="#">MandrakeSecure</a> <a href="#">MandrakeOnline</a>
Mailing Lists	<a href="#">Mandrake lists</a> <a href="#">Archives - newbie</a> <a href="#">Archives - expert</a> <a href="#">Archives - cooker</a>
Download Links	<a href="#">Download mirrors</a>
Documentation	<a href="#">Mandrake 9.0 Documentation</a>
External Resources	<a href="#">DistroWatch: Mandrake page</a> <a href="#">Desktop Mandrake</a> <a href="#">Mandrake Users Board</a> <a href="#">Penguin Liberation Front lair</a> <a href="#">Texs tar's RPMs</a> <a href="#">Buchan Milne's RPMs</a>

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

URL: <http://www.distrowatch.com/dwres.php?resource=review-mandrake>

tenets of computer security.

Microsoft, as the US trial record and experience has shown, has added these complex chunks of code to its operating system not because such programming complexity is necessary, but because it all but guarantees that computer makers, users and consumers will use Microsoft products rather than a competitor's.

These competition related security problems have been with us, and getting worse, for years. The recent spate of virus attacks on the Internet is one more sign that we must realize the danger we are in. The report CyberInsecurity - The Cost of Monopoly is a wake up call that government and industry need to hear.

September 24, 2003

## **CYBER/INSECURITY: THE COST OF MONOPOLY HOW THE DOMINANCE OF MICROSOFT'S PRODUCTS POSES A RISK TO SECURITY**

### **Executive Summary**

Computing is crucial to the infrastructure of advanced countries. Yet, as fast as the world's computing infrastructure is growing, security vulnerabilities within it are growing faster still. The security situation is deteriorating, and that deterioration compounds when nearly all computers in the hands of end users rely on a single operating system subject to the same vulnerabilities the world over.

Most of the world's computers run Microsoft's operating systems, thus most of the world's computers are vulnerable to the same viruses and worms at the same time. The only way to stop this is to avoid monoculture in computer operating systems, and for reasons just as reasonable and obvious as avoiding monoculture in farming. Microsoft exacerbates this problem via a wide range of practices that lock users to its platform. The impact on security of this lock-in is real and endangers society.

Because Microsoft's near-monopoly status itself magnifies security risk, it is essential that society become less dependent on a single operating system from a single vendor if our critical infrastructure is not to be disrupted in a single blow. The goal must be to break the monoculture. Efforts by Microsoft to improve security will fail if their side effect is to increase user-level lock-in. Microsoft must not be allowed to impose new restrictions on its customers - imposed in the way only a monopoly can do - and then claim that such exercise of monopoly power is somehow a solution to the security problems inherent in its products. The prevalence of security flaw in Microsoft's products is an effect of monopoly power; it must not be allowed to become a reinforcer.

Governments must set an example with their own internal policies and with the regulations they impose on industries critical to their societies. They must confront the security effects of monopoly and acknowledge that competition policy is entangled with security policy from this point forward.

---

*The threats to international security posed by Windows are significant, and must be addressed quickly. We discuss here in turn the problem in principle, Microsoft and its actions in relation to those principles, and the social and economic implications for risk management and policy. The points to be made are enumerated at the outset of each section, and then discussed.*

## **1. THE PROBLEM IN PRINCIPLE**

*To sum up this section:*

- Our society's infrastructure can no longer function without computers and networks.
- The sum of the world's networked computers is a rapidly increasing force multiplier.
- A monoculture of networked computers is a convenient and susceptible reservoir of platforms from which to launch attacks; these attacks can and do cascade.
- This susceptibility cannot be mitigated without addressing the issue of that monoculture.
- Risk diversification is a primary defense against aggregated risk when that risk cannot otherwise be addressed; monocultures create aggregated risk like nothing else.
- The growth in risk is chiefly amongst unsophisticated users and is accelerating.
- Uncorrected market failures can create and perpetuate societal threat; the existence of societal threat may indicate the need for corrective intervention.

### Discussion

Computing is essential to industrialized societies. As time passes, all societal functions become more deeply dependent on it: power infrastructure, food distribution, air traffic control, emergency services, banking, telecommunications, and virtually every other large scale endeavor is today coordinated and controlled by networked computers. Attacking national infrastructures is also done with computers - often hijacked computers. Thus, threats to computing infrastructures are explicitly and inherently risk harm to those very societies in proportion to those society's dependence on them. A prior history of catastrophe is not required to make such a finding. You should not have to wait until people die to address risks of the scale and scope discussed here.

Regardless of where or how it is used, computing increases the capabilities and the power of those who use it. Using strategic or military terminology that means what it sounds like computing is a "force multiplier" to those who use them - it magnifies their power, for good or ill. The best estimates of the number

of network connected computers show an increase of 50% per year on a worldwide basis. By most general measures what you can buy for the same amount of money doubles every eighteen months ("Moore's Law"). With a conservative estimate of a four year lifetime for a computer - in other words, consumers replace computers every four years on average - the total computing power on the Internet therefore increases by a factor of 2.7 per annum (or doubles every 10 months). If a constant fraction of computers are under threat of misuse, then the force available to misusers will thus double every 10 months. In other words, the power available to misusers - computer hackers in popular parlance - is rising both because what they can buy grows in power per dollar spent and because the total number of networked computers grows, too. Note also that this analysis does not even include attacks enabled by storage capacity, which doubles in price-performance twice as fast as CPU (doubles every nine months rather than eighteen).

Internetworked computing power makes communication feasible. Communication is of such high value that it has been the focus of much study and much conjecture and not just recently. For one-way broadcast communication, the value of the network itself rises proportionally to  $N$ , the potential number of listeners ("Sarnoff's Law"). By way of example, advertisers pay for television time in rough proportion to the number of people viewing a given program.

For two-way interactive communications - such as between fax machines or personal email - the value of the network rises proportionally to  $N^2$ , the square of the potential number of users ("Metcalfe's Law"). Thus, if the number of people on email doubles in a given year, the number of possible communications rises by a factor of four.

Growth in communications rises even more when people can organize in groups, so that any random group of people can communicate with another. Web pages, electronic mailing lists and online newsgroups are good examples of such communications. In these cases, the value of the network rises proportionally to  $2^N$ , the potential number of groups being an exponential growth in  $N$  ("Reed's Law").

Assume for now that the Internet is somewhere between the Metcalfe model, where communications vary according to the square of the number of participants ( $N^2$ ), and the Reed model, where communications vary according to two raised to the  $N$ th power ( $2^N$ ).

If we make this assumption, then the potential value of communications that the Internet enables will rise somewhere between  $1.5^2 = 2.3$  and  $2^{1.5} = 2.8$  times per annum. These laws are likely not precisely accurate. Nonetheless, their wide acceptance and historic record show that they are good indicators of the importance of communication technology.

To extend this simple mathematical model one final step, we have assumed so far that all communications

are good, and assigned to the value of the network a positive number. Nonetheless, it is obvious that not all communications (over computer networks, at least) are positive. Hackers, crackers, terrorists and garden-variety criminals use the network to defraud, spy and generally wreak havoc on a continual basis. To these communications we assign a negative value.

The fraction of communications that has positive value is one crucial measure, and the absolute number of negative communications is another. Both are dependent on the number of networked devices in total. This growth in the number of networked devices, however, is almost entirely at the "edges" of networked computing - the desktop, the workstation, the home, the embedded system, the automated apparatus. In other words, the growth in " $N$ " is not in the core infrastructure of the Internet where highly trained specialists watch over costly equipment with an eye towards preventing and responding to attacks. Growth, rather, is occurring mostly among ordinary consumers and non-technical personnel who are the most vulnerable to illegal intrusions, viruses, Trojan horse programs and the like. This growth at the periphery, furthermore, is accelerating as mobile, wireless devices come into their own and bring with them still more vulnerabilities.

Viruses, worms, Trojan horses and the like permit malicious attackers to seize control of large numbers of computers at the edge of the network. Malicious attackers do not, in other words, have to invest in these computers themselves - they have only to exploit the vulnerabilities in other people's investments.

Barring such physical events as 9/11, an attack on computing is a set of communications that take advantage of latent flaws already then present in those computers' software. Given enough knowledge of how a piece of software works, an attacker can force it to do things for which it was never designed. Such abuse can take many forms; a naturalist would say that attacks are a broad genus with many species. Within this genus of attacks, species include everything from denial of service, to escalation of authority, to diversion of funds or data, and on. As in nature, some species are more common than others.

Similarly, not all attacks are created equal. An annoying message that pops up once a year on screen to tell a computer user that he has been infected by Virus XYZ is no more than that; an annoyance. Other exploitations cost society many, many dollars in lost data, lost productivity and projects destroyed from data crashes. Examples are many and familiar including the well known ILOVE YOU, NIMDA, and Slammer attacks not to mention taking over users machines for spamming, porn distribution, and so forth. Still other vulnerabilities, though exploited every day and costing society substantial sums of time and money, seldom appear in the popular press. According to London-based computer security firm, mi2g Ltd., global damage from malicious software inflicted as much as \$107 billion in global economic damage this year. It estimates that the SoBig worm, which helped make August the

costliest month in terms of economic damage, was responsible for nearly \$30 billion in damage alone.[1]

For an attack to be a genuine societal-scale threat, either the target must be unique and indispensable - a military or government computer, authoritative time lookup, the computer handling emergency response (911) calls, airport flight control, say - or the attack must be one which once triggered uncontrollably cascades from one machine to the next. The NIMDA and Slammer worms that attacked millions of Windows-based computers were examples of such "cascade failure" - they spread from one to another computer at high rates. Why? Because these worms did not have to guess much about the target computers because nearly all computers have the same vulnerabilities.

Unique, valuable targets are identifiable so we, as a society, can concentrate force around them. Given enough people and training (a tall order to be sure), it is possible to protect the unique and core assets. Advanced societies have largely made these investments, and unmitigated failures do not generally occur in these systems.

Not so outside this core: As a practical and perhaps obvious fact, the risk of cascade failure rises at the edges of the network where end users are far more likely to be deceived by a clever virus writer or a random intruder. To put the problem in military terms, we are the most vulnerable when the ratio of available operational skill to available force multiplication is minimized and thus effective control is weakest. Low available skill coupled to high potential force multiplication is a fair description of what is today accumulating on the periphery of the computing infrastructures of every advanced nation. In plainer terms, the power on the average desktop goes up very fast while the spread of computers to new places ensures the average skill of the user goes down. The average user is not, does not want to be, and should not need to be a computer security expert any more than an airplane passenger wants to or should need[1] to be an expert in aerodynamics or piloting. This very lack of sophisticated end users renders our society at risk to a threat that is becoming more prevalent and more sophisticated.

Regardless of the topic - computing versus electric power generation versus air defense - survivability is all about preparing for failure so as to survive it. Survivability, whether as a concept or as a measure, is built on two pillars: replicated provisioning and diversified risk. Replicated ("redundant") provisioning ensures that any entity's activities can be duplicated by some other activity; high availability database systems are such an example in computing just as backup generators are in electric power. The ability of redundant systems to protect against random faults is cost effective and well documented.

By contrast, redundancy has little ability to protect against cascade failure; having more computers with the same vulnerabilities cannot help if an attack can

reach them all. Protection from cascade failure is instead the province of risk diversification - that is, using more than one kind of computer or device, more than one brand of operating system, which in turns assures that attacks will be limited in their effectiveness. This fundamental principle assures that, like farmers who grow more than one crop, those of us who depend on computers will not see them all fail when the next blight hits. This sort of diversification is widely accepted in almost every sector of society from finance to agriculture to telecommunications. In the broadest sense, economic diversification is as much the hallmark of free societies as monopoly is the hallmark of central planning.

Governments in free market societies have intervened in market failures - preemptively where failure was be intolerable and responsively when failure had become self-evident. In free market economies as in life, some failure is essential; the "creative destruction" of markets builds more than it breaks. Wise governments are those able to distinguish that which must be tolerated as it cannot be changed from that which must be changed as it cannot be tolerated. The reapportionment of risk and responsibility through regulatory intervention embodies that wisdom in action. If governments are going to be responsible for the survivability of our technological infrastructure, then whatever governments do will have to take Microsoft's dominance into consideration.

---

Next installment will cover Microsoft and the impact on public protection.

[1] "Government Issue", David Zeiler, The Baltimore Sun/SunSpot.net. September 18, 2003.

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

<http://www.cciinet.org/papers/cyberinsecurity.pdf>

## An Overview of Sentinel3G

Authors: Alan Main <[alan.main@fs.com.au](mailto:alan.main@fs.com.au)>

Most commercial data centres provide services that are critical to the day-to-day functioning of their entire organisation and everybody agrees these services and the infrastructure supporting them (hardware, operating system, database etc) should be monitored to ensure that everything is "healthy".

The simplest form of monitoring is to write some scripts (e.g. ping a remote host) that are periodically run via CRON. When these scripts finds something unexpected (such as the host not responding to the PING or that the application is down), the script notifies someone (perhaps via email or an SMS message) that there is a problem. The problem with home grown scripts is this

can take a lot of time to write and maintain, especially if the environment changes frequently or the scripts are to be used by other people.

While there are a number of commercial and free products available, we found that people were generally not particularly satisfied with them for a number of reasons:

- The low cost and free products are often no simpler to maintain or more flexible to use than the home grown script and many of high cost products aren't much better!
- Those few with the features that are required by most medium enterprise are both too expensive and too complicated to configure, and tend to be heavy on system resources.
- Many are oriented towards network monitoring, rather than generic system monitoring.
- Most lack the flexibility to be able to "plug in" existing scripts written in any language as agents.

As a result we believed there is a market for a low-cost, flexible, network-wide System Monitoring product aimed at small to medium sized enterprises with good reporting and response capabilities. We also felt that many large organisations would also be interested in avoiding the expense and complexity of existing "high-end" products. This has indeed proved to be the case. This led us to develop Sentinel3G, a flexible and powerful tool for system monitoring. It provides a structured management environment in which define rules about the things being monitored and how to respond to event as they occur. For such a powerful product it is simple to use, and has an extremely light footprint.

If you go to our web site, [www.sentinel3g.com](http://www.sentinel3g.com), you can look at the live demonstration or download a free version of Sentinel Netservices Lite for Linux.

## WHY USE SENTINEL3G

Download install & start using it in less than 15 min  
sentinel3G is easy-to-use, simple to customize and quick to deploy. With its small download size (approximately 4.5MB) and a simple one-step, self-configuring installation, you can have sentinel3G up and running on your system in less than 15minutes.  
Network-wide systems monitoring from one place

sentinel3G enables you to see the state of all your systems from one convenient place. Each machine runs a Host Monitor that uses agents to gather data about the state of the services on that system, e.g. CPU, memory, filesystem, web server, database. All this information, from all your systems, can be displayed on a single console, allowing an overview of your entire network at a glance. sentinel3G can even send alerts to your mail, pager or mobile phone.

## RAPID DIAGNOSTICS

Console views organize information so that even the largest networks can be understood at a glance. You can quickly diagnose problems by drilling down to find

the state of any service. sentinel3G consoles are available for web browsers, Microsoft Windows or X Windows so you can monitor your systems from anywhere using almost anything.

## PROACTIVE MANAGEMENT

sentinel3G goes beyond simple passive monitoring. When a service changes state it can automatically run a series of responses to proactively manage the service. For many routine situations these automated responses can resolve the problem, saving you time and effort.

## SMART AND SIMPLE

Every service is monitored by a "Sentry" which has a set of user-defined states. These states, defined by powerful but easy to create rules, describe the status of the service in a meaningful way, e.g. overloaded, busy, working and idle. As a Sentry changes state its console icon is changed to reflect that state.

## MONITOR ANYTHING

sentinel3G can monitor just about anything. Any service that provides data can have a sentinel3G agent configured to collect that information. Through simple configuration screens you can quickly create a Sentry with the states, responses and icons necessary to monitor and automatically manage that service.

## ADD ON KNOWLEDGE BASES

The free version of sentinel3G for Linux comes with a standard Knowledge Base (KB), NetServices Lite, for monitoring the availability of networked systems and a Linux Lite KB to monitor the basics of the Linux OS. Additional KBs can be purchased for other things such as Solaris, HP-UX, AIX, Linux, Windows NT/2000™, Oracle, Informix, Apache, Squid etc. You can also quickly and easily develop your own site-specific KBs to meet your local requirements.

## REPORTING, GRAPHING AND LOGGING

sentinel3G supports real-time and historical graphing and reporting on information collected about your services. The DataLogging KB keeps data, event and audit logs which can be maintained for as long as you wish allowing problems to be tracked and diagnosed long after the event. The EventHistory KB allows detailed service level reports to be created, complete with charts and graphs.

## SIMPLIFIED CONFIGURATION

Configuration is managed via a graphical user interface eliminating complicated manual editing of configuration files. With a little experience new Sentries and Agents can be quickly added to monitor and manage any of the services on your systems.

## A STATE-MACHINE ARCHITECTURE

The essence of sentinel3G is a finite-state machine. Each thing to be monitored is represented as a sentry that can be in one of an arbitrary set of states, with each state having an associated severity. A sentry has a set of conditions that define each state. These

conditions are evaluated based on data returned from one or more agents. An event is created when data from an agent causes a sentry to change states.

This table shows a simple example of the states of a disk space sentry. The idea is that the state name would convey something meaningful. Notification of problems, execution of responses, reports and documentation to assist users diagnose problems are associated with the appropriate state.

State	Severity	Entry Condition	Delay	Responses
Delete	-	no-data		
Full	critical	\$pct_free == 0	0	Notify the boss
Very low	Alarm	\$pct_free < 2	360	Notify admin
Low	warning	\$pct_free < 10	300	Remove temp & core files
OK	normal	-	300	Notify admin

Each state can have a sequence of responses attached to it each with an optional delay before each is performed. These responses can be arranged in an escalation pattern so if the first fails the next is tried or if the state persists a more senior person is notified. For example, "admin" will be notified immediate the Very Low state is entered and the boss is notified if the state persists for 3600 seconds.

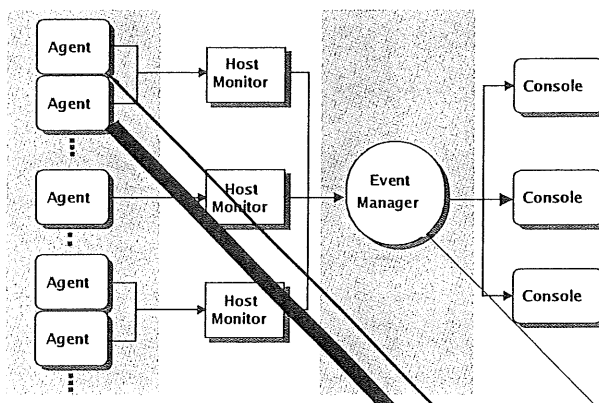
Our design requires that Entry Conditions be able to manipulate numeric and string variables and specify arbitrary boolean conditions so we decided to leverage an existing language (Tcl) rather than defining our own and building a parser. This means state changes can be built from very complex and powerful expressions.

### THE SYSTEM ARCHITECTURE IN BRIEF

sentinel3G was designed to monitor a small or large number of distributed systems (servers, applications, routers etc) on a network and is build using an appropriate light-weight distributed architecture that is described below.

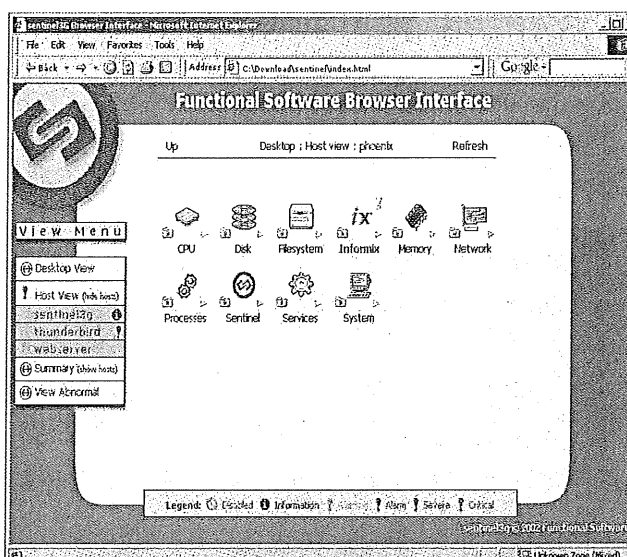
The Host Monitor is a per-host daemon process that receives data from its agents, detects events (see A State-Machine Architecture below?), performs data logging and notifies the central Event Manager when events occur. It performs discovery and all responses, including notification, and is responsible for scheduling the agents.

The Host Monitor is the main "engine" of sentinel3G and each Host Monitor runs independently, ensuring that its host will be monitored regardless of whether other hosts (including the Event Manager host) or networks are down.



The Event Manager is the central process that collects state information from all Host Monitors, and updates the data on the Consoles when required. Only state change and updated console information is sent across the network from each Host Monitor to the central Event Manager, keeping network traffic to a minimum.

The Console is the primary user interface to sentinel3G, providing one or more hierarchical views of the sentries being monitored. Each sentry is displayed as an icon with text. The Console allows arbitrary hierarchical grouping of sentries in folders. When an event occurs, the corresponding sentry changes state, e.g. a printer goes from "up" to "down". The user can "drill-down" the hierarchy to get more details, or display different "views", for example, to display only abnormal sentries of a particular type. Access security may also be assigned to folders providing tailored views of certain sentries to certain groups of users.



The flow of data through the system is as follows:

- An agent returns some data to the Host Monitor, setting some variables;
- The state conditions of all sentries using the agent are then evaluated in order. The first to evaluate to "true" causes a change to that state;
- If a state change occurs the Host Monitor:  
  - informs Event Manager (if it is up);

- logs the state change;
- cancels any pending responses for the previous state;
- schedules any associated responses for the new state;
- Any changes to the console text, icons etc are sent to the Event Manager (if it is up), and from there to the active consoles;

## IMPLEMENTATION

Our previous years of experience (> 15) in developing complex systems tools taught us that we could build sentinel3G far more effectively and efficiently in an interpretive language with a small component in "C". This choice made sentinel3G very compact, highly robust. It has enabled us to build it very quickly as interpretive languages tend to make developing, testing and debugging very rapid compared to traditional approaches based on "C". Further more, and surprisingly to many, we have found that systems tools built in this way are more efficient, due to a greater focus on architecture and algorithms, rather than coding and debugging.

We chose Tcl (with incr Tcl object-oriented extensions) over other languages such as Java or C++ because:

- A subset of Tcl syntax was used in Sentinel3G's own configuration (state conditions, variable expressions etc). Tcl is able to directly parse and execute these at runtime.
- Tcl can be used as an extension language both by end-users to develop agents.
- Tcl has a number of neat extensions such as scotty, for network & SNMP support.
- Tcl is a mature, stable Open Source product, and a number of different groups are working on it and extensions.

*Although sentinel3G is not Open Source, all the Tcl source code is provided, which makes it feasible for users to add new agent interfaces and data manipulation functions.*

## Supercluster Redux

Frank Crawford (FC) & Vladas Leonas (VL) discuss Supercomputer Clusters

FC: In 2002 five NSW universities who are part of the ac3 University consortium put in a funding submission to the Australian Research Council.

ac3 (the company) wasn't involved at that stage. At the time the Universities did not want a 64-bit architecture, a cluster of machines, probably based on Linux, to address the needs of people doing multi-processing to resolve several classes of problems.

At that time they were considering a loosely coupled system; they didn't want a shared memory machine, as there was another shared memory system already deployed at ac3.

When did ac3 start to get involved?

VL: We got involved when they got a word that the grant submission was successful.

What was the value of the grant?

FC: The budget they had was around \$780,000 - about half contributed by the Universities and the rest was an ARC grant. The idea was being to spend \$740,000 on the machine and about \$40,000 for the software.

What happened then?

VL: ac3 began to act as a strong enabler. Once they had the grant we sat down with the universities and extracted their requirements.

A few things changed between the submission and what we extracted from them. For example, they decided they wanted 32-bit architecture rather than a 64-bit architecture considered at some stage. The reason being, the kinds of problems they wanted resolved were well within the capability of 32-bit architecture.

Then we asked about their requirements for the interconnect - for interprocess communications. The question was, did they want a separate network for storage or did they want a single network for everything?

The size of secondary storage was an issue too: they said they needed about a terabyte and a half of RAID 5 useable disk space.

Once we'd agreed on those high level requirements we moved onto more low-level requirements around 32-bit architecture.

The university researchers were pretty strong on the speed requirements of the front side bus. They said a 400 MHz FSB wouldn't be fast enough. They really wanted a 533 MHz FSB, which was the fastest at the time.

They insisted on genuine Intel CPUs than AMD because they intended to use Intel compilers that were heavily optimised against Intel CPUs.

They put certain benchmarks that each node and interconnect had to comply with. So we finally ended up with a tender document that actually specified their requirements in very "black and white" terms.

When did it go to tender?

FC: December 2002. Ten companies responded to the tender. There were three big tier one companies, two tier two companies and the rest were tier three companies. Initially we shortlisted three submissions for further discussion.

The tender document required demonstrable experience in building clusters and project management of building clusters. We didn't want a DIY approach; we



asked the successful tenderer to build and commission it on site.

We planned to create a production box with high level of availability and a three-year on site warranty.

When did you appoint the successful tender?

VL: The whole process took six months, which was a bit longer than we would have expected. The initial phase was short – we were able to narrow down the field of acceptable candidates quite quickly.

After that we had a longer phase of detailed discussions with prospective companies, and then a period of nitty gritty discussions with Dell.

In the end, Dell got the tick because the other tenderers couldn't match their overall response.

FC: One unexpected bonus of the delay was that the performance and pricing of the systems improved markedly – at one stage I said if we delayed this any longer we'll have the fastest machine ever built!

The original tender document called for 2.4GHz CPUs (hoping to get 2.8) but during the intervening period leading up to the signing of the contract we ended up getting 3.06GHz CPUs for a lesser sum of money.

VL: Another point to note was the tender requirement for scalability. We wanted the cluster to be able to grow, especially in the first year of operation. There were two aspects to this.

The first was that the switch had to allow more nodes to be connected and simultaneously to maintain certain levels of redundancy. The second was about fixing the price for the first 12 months if we chose to grow the cluster.

How much can it grow by?

VL: We can probably add another 32 to 48 nodes before we start hitting certain issues on the switch.

And what will that mean for performance at that point?

FC: The answer to that is scaled linearly. We've got 155 nodes at present. Adding another 32 to 48 nodes would add about 25 per cent to the machine's performance.

That is, we can have 25 per cent more jobs running simultaneously or individual users would get a faster turnaround by virtue of having more CPUs devoted to their job.

What are the machine's specifications?

FC: Once the machine was put together the issue of benchmarking became important. While there were benchmark criteria in the contract, one of the big issues for the university people was to get a machine in Top500 and find where they fitted in there.

There was also a hope at that stage to make it the fastest machine in Australia although that wasn't an initial goal of the project. What kicked this ambition along was the fact that we had additional research partners come in after the contract was let who bought an additional 24 nodes.

So we went through the Top500 Linpack benchmarking process (<http://www.top500.org/lists/linpack.php>). The important point to note here is the difference between theoretical and sustained peak performance. The machine's theoretical peak performance is 1.87 Teraflops.

The more important figure is sustained performance. Our machine's sustained performance on the Linpack industry standard benchmark is 1.095 Teraflops, which is about 60 per cent efficiency. Sixty per cent efficiency, I might add, is about the best you can get for cluster machines anywhere in the world.

Another point to note is performance at the individual node level because that's what makes a difference to individual users. Our nodes are dual Pentium 3.06 GHz CPUs that are significantly faster than the 2.4 GHz nodes that were common only six months ago.

On large jobs the thing that affects performance is the network behind the scene and this is where we selected a Gigabit Ethernet switch.

VL: The choice of a Gigabit Ethernet switch makes the cluster as a whole the best on price performance.

The reason this is important on larger jobs is that we want the nodes to talk to each other quickly, right?

FC: There are two issues here. There's the latency, which is basically how long it takes when you "start talking" and when the other end starts to "hear" you, and then there's bandwidth, which is how big a message you can transmit.

In the gigabit switch we have it actually has a very low latency and a large bandwidth. Most of the other cluster machines around at the moment are using fast Ethernet switches or specialised network connections. The fast Ethernet switch have higher latency and have lower bandwidth because they're only 100-megabit connections, which the specialised networks connections are much more expensive.

We could have gone with even faster switches but in the end we would have only got 15 per cent more efficiency for about double the price. In the end, we couldn't really justify that kind of investment for so little gain.

What about the Front-Side Bus?

FC: Because many of the jobs are fairly memory-intensive one of the big issues is access to the fast memory in the machine.

Eighteen months ago the fastest you'd get was 166 MHz

front-side-Bus. Front side Bus is an Intel term that describes the speed at which a CPU communicates with RAM memory. Recently performance has improved with the newer 533 MHz FSB, which is what we've got on the machine.

The bottom line is we get faster performance at the individual node level for jobs that require RAM access.

VL: It's worth noting that we have tried to optimise the machine at two levels. That is, at the individual node level and at the whole configuration level.

Tell us the Linux story.

FC: Just about anyone who's building a cluster machine these days is building a Linux cluster. The Linux system we're using is based on Red Hat 9. This is important for our users because they are picking up code from a number of people around the world that will pretty much run on the cluster.

VL: Researchers share certain applications that are run on clusters and the Linux environment facilitates this. A lot of the codes that are running on Linux, particularly the applications for high performance computing, were brought across from various Unix systems that have been going for a number of years, so there's a lot of inherited code.

FC: Another important point about the open source environment is that we were able to go and use a number of open source packages.

We're using a package called OSCAR – it stands for Open Source Cluster Application Resource – which is designed to make it easy to put together clusters. To some extent OSCAR is a meta-package, since it brings together a number of other open source packages in a form that is useful for clusters.

Tell me about the official launch.

VL: ac3's new Dell-Linux supercomputer was officially launched on the 25<sup>th</sup> November by NSW Minister for Commerce, the Hon John Della Bosca, MLC and Prof Lesley Johnson of UTS.

The Minister gave a great speech to around 80 guests from the business, government and research sectors. He noted that the machine was being used for research projects in photonic technology, drug design, mobile telephony and the modelling land and environmental changes.

He also noted that on international rankings the machine is in the top 20 per cent of fastest supercomputers in the world.

VL: We haven't formally released the cluster to all our users at this point. We are still in the so-called preview-users mode, when selected users have been allowed to start using the machine.

The expected number of registered users will grow from

about 20 to 30 now to around 200-300 over time. Even with this small number of current users, the machine is being heavily used.

At any given point, half the machine is being used at 75 per cent capacity and another five to 10 per cent is used at 100 per cent capacity. So average utilisation is currently running at 75 to 80 per cent. This figure is increasing all the time.

## Linux Desktop Underrated

Author: Tom Adelstein <adelste@netscape.net>

Many writers have suggested that the first ever Desktop Linux Consortium conference, held November 10th near Boston, indicated a sign of the acceptance of Linux in the enterprise. As one of the speakers, I took my invitation as an opportunity to analyze the event. I walked away convinced that many of the speakers lacked the knowledge, will and/or ability to convey the real status of the Linux desktop. I felt they undersold the Linux Desktop. I note the exceptions later in the article.

In dealing with the underselling of the Linux desktop by its proponents, I intend to show that desktop Linux can help enterprises now, especially those short of resources like city governments.

The Highlight Speakers

I anticipated that the IBM presentation (<http://www.desktoplinux.com/news/NS9189274301.htm>) would be well attended since it was the only event scheduled in its time slot. Having significant knowledge of IBM's desktop status, I was disappointed by their speaker. Sam Docknevich's presentation (<http://www.desktoplinux.com/news/NS9189274301.htm>) provided clear evidence of the blind leading the blind and IBM's sluggishness in committing to open desktop offerings. Docknevich runs the IBM Global Services grid initiatives, so what was he doing making the Desktop presentation?

I kept asking myself what credentials Sam had to discuss the desktop. So it was no surprise that when he took questions from the audience, he was lost. He represented the high-tech equivalent of a politician. Only, most of Docknevich's answers started and ended with "I don't know". Politicians rarely say, "I don't know". They prefer the taste of shoe leather, I guess. Docknevich made a mad dash to the exit after his presentation. I don't blame him.

Nat Friedman of Ximian (<http://ximian.com/about-us/>) gave the most compelling and enlightening perspective on the desktop. What Nat failed to tell the audience still remains untold. Ximian could easily have every bell and whistle to surpass both the Apple Mac OS X (<http://www.apple.com/macosx/>) desktop and Microsoft XP (<http://www.microsoft.com/windowsxp/default.aspXP>). The only reason they don't own the desktop space

relates to underfunding. That began with their original Venture Capitalists (<http://www.crv.com/>) and extends to Novell (<http://novell.com/>). If Novell gives these guys the resources they need, Linux will ramp up faster than Mac OS X did.

Which leads us to the absence of the people with the real Linux Desktop: Sun Microsystems. Sun's Jonathan Schwartz ([http://www.sun.com/aboutsun/media/ceo/mgt\\_schwartz.html](http://www.sun.com/aboutsun/media/ceo/mgt_schwartz.html)) seemed conspicuous by his absence. In my three week evaluation of several Linux desktops, I concluded that Sun's new Linux desktop outpaced everyone. Sun did not make themselves available at the conference.

The irony of Sun's absence has much to do with Novell's recent purchases of both Ximian and SuSE. Sun contracted with SuSE to build the underlying Linux operating system for the Java Desktop. Ximian built the first Gnome desktop for Sun. Sun Microsystems originally sponsored the Gnome Foundation (<http://foundation.gnome.org/>) which Ximian uses as the core of their desktop. Sun should have been well represented and their absence was a glaring oversight by conference planners. With luck this will be remedied next year, after Sun puts a few more China's (<http://www.sun.com/smi/Press/sunflash/2003-11/sunflash.20031117.3.html>) up on the board to emphatically stamp its Linux Desktop credibility and leadership in everyone's minds.

## SUN'S STRATEGY: SOLID

The Java Desktop (<http://www.sun.com/software/javadesktopsystem/index.html>) provides the best user experience of any on the market. I prefer it to Mac OS X Panther and Microsoft's XP. When users begin to experiment with Sun's desktop, I expect the WOW effect to take place.

Anyone in the broad IT field knows that Sun's engineering ([http://www.sun.com/corp\\_emp/zone/special.html#engr](http://www.sun.com/corp_emp/zone/special.html#engr)) surpasses everyone else's. Sun put that same quality into their Java Desktop that one has come to expect from anything Sun. The user gets the benefits of Sun's cumulative talents. Expect an incredible desktop from Scott McNealy's ([http://www.sun.com/aboutsun/media/ceo/mgt\\_mcnealy.html](http://www.sun.com/aboutsun/media/ceo/mgt_mcnealy.html)) band of merry engineers.

Sun's desktop also will surprise the world when it comes time to market it. Sun has what the general business world doesn't have: a clear target into the Microsoft user base.

Unlike Sun, you will find little, if any media information about the current Microsoft user base. From my research, approximately 400 million Windows 98 desktops remain in the installed base. Some estimates indicate that Windows XP has an installed base of 35 million in comparison. If these figures resemble the

market, then Sun has a chance to take the desktop market.

Compatibility presents a big problems within the so-called homogeneous Windows installed base (<http://www.microsoft.com/windows/WinHistoryDesktop.msp>). For example, users with Office XP or 2000 do not have the ability to collaborate with users of earlier versions of Office. So, Sun provides a free office productivity suite that allows all Office users to collaborate.

You also have the problem of old Windows clients being able to work with Microsoft's active directory (<http://www.microsoft.com/windowsserver2003/technet/overview/activedirectory.msp>). Sun doesn't have that problem. Sun's Linux authenticates with older Win32 server platforms and with the new generation of Windows clients and servers. No one will have to buy new hardware to run Linux unlike the Microsoft offerings.

## SUN, OPENOFFICE AND STAROFFICE 7

I have observed corporate, government and academic users adding OpenOffice (<http://openoffice.org/>) to their systems just to read the various formats Microsoft produces. Many of those users begin migrating to OpenOffice as the default productivity suite.

Now, Sun can follow up with users who may already have become accustomed to OpenOffice. If that's the case, no migration issues will exist at the user level. In fact, the StarOffice 7 (<http://www.sun.com/software/star/staroffice/features.html>) suite provides a quality experience even compared with the newest Microsoft products. I actually prefer StarOffice 7.

If Sun will enhance some rendering features of OpenOffice, they will find more users migrating, which puts them in a front-row seat for a painless migration to StarOffice 7, as options warrant. Additionally, Sun provides a solution to the so-called homogeneous Microsoft environment which in reality looks more like a heterogeneous mess. Adding OpenOffice or StarOffice 7 allows a Windows 98 user to read documents written in the gambit of Microsoft Word 6 through XP.

## THE OTHER LINUX DESKTOPS

Prior to the Desktop Linux Conference, I did extensive testing and evaluation of several Linux distributions including Sun's Live CD Beta, Lycoris, Xandros 1.1, SuSE 9.0, Red Hat 9.0, Gentoo, Debian and Ximian Desktop 2 on SuSE 8.2 and Red Hat 9. I kept my home directory on a CD-RW and updated as needed.

I learned how to tune and enhance the Linux desktop from Nicholas Donovan back in 1999. Since then, I have progressed in making the desktop work. With CodeWeavers CrossOver Office (<http://codeweavers.com/>), I found myself able to ditch the last of my Windows distributions (<http://microsoft.com/windows>) this year.

## LYCORIS

Lycoris (<http://www.lycoris.com/>) provides one of my favorite GUI look&feels. In fact, I learned some lessons on creating a pleasant experience from Lycoris. If Red Hat's Matthew Szulik (<http://www.redhat.com/about/corporate/team/szulik.html>) thinks Windows provides a good OS for his 90 year-old father, he missed the boat by not checking out Lycoris.

While Lycoris downplays their acquisition of Caldera Linux (<http://sco.com/>) as the core of their distribution, some users may want to consider this in a purchase decision. Lycoris, like Sun Microsystems, owns the System V UNIX license. That takes some of the cognitive dissonance out of your purchase.

## XANDROS 1.1

I didn't look at the Xandros 2 distribution since it wasn't ready at review time. I did look at their Technology Preview. I understand that version 2.0 uses KDE 3.x (<http://kde.org/>) instead of KDE 2.x used in versions 1.0 and 1.1.

Xandros (<http://xandros.com/>) has some remarkable technology. The highlights of their product include it's ability to resize and partition the disk without disturbing an existing Windows installation. That feature should belong in every distributions of Linux.

Xandros also has a "network neighborhood" style file manager that maps NFS and Windows shares. That feature began in Corel Linux and still amazes me. Consider Xandros interoperability the best of the rest. Sun also has similar and excellent interoperability.

Xandros uses Debian as the underlying Linux distribution. Most people consider the Debian apt-get update service the best in the Linux community. Xandros has taken that a few steps further and added a very nice graphic user interface. Rank Xandros number one in the update/upgrade category until we see the final Sun product.

Xandros also puts a high quality control spin on the Debian packages. The developers actually fix the original Debian packages if even slightly broken. So, the applications from the Xandros Network has the highest quality of the various packages.

Aside from the nice polished look and feel, Xandros includes Codeweavers Crossover products. So, one gets a ready to rock and roll distribution. Xandros appears to have the best driver support of all Linux distributions. I'll bet that even Matthew Szulik's father could make this work.

## SUSE 9.0

I think Jack Messman (<http://www.novell.com/company/bios/ceo.html>) blessed the Linux community by purchasing SuSE(<http://suse.com/>) The reputation for high quality German engineering never made it into this company. If you have a penchant for mind reading, you

could make this product work for you. I'm sure that once Nat and Miguel get their hands on SuSE, you'll see a product to rival the Sun desktop. Until then, buy it as a server and use Novell Nterprise additions. Then you have a great server product with the administrator in mind.

As a desktop, SuSE is a good product for experienced Linux users who want to spend a lot of money for a commercial distribution. The only problem you might see lies in the fact that experienced Linux users wouldn't buy a product like SuSE for their desktop. I bought it to test and would like a refund. I did eventually tweak KDE by adding lots of icons, web browsers and simple development tools.

If someone looked at SuSE and judged Linux accordingly, they would conclude Linux isn't ready for the desktop. So do not blame every journalist for making that mistake. Blame IBM for using SuSE exclusively and then concluding Linux isn't ready for the desktop.

## RED HAT 9.0

Can you say "good riddance". The hardest Linux company with whom to deal has taken its branding away and rode off into the Sunset. Frankly, I will miss some things about them. I won't miss their corporate culture.

I used Red Hat (<http://redhat.com/>) as my desktop for the last few years starting with version 5.0. Two features worked for me. They did a good job a rendering the desktop and they had up2date. Somehow, I felt comfortable seeing that up2date icon on my desktop turn from red to blue when I did my updates.

You can buy their Professional Workstation if you want to continue with Red Hat. After my entitlement runs out in April, I'll have moved to another distribution. Fedora Core 1 has no place in my future.

My Red Hat Linux desktop hums nicely. It does everything I want. However, I have to tweak it and tweak it and tweak it. I use KDE instead of the default Gnome desktop. I also have to find applications to add. While I found nothing easy about Red Hat, I got used to it.

Like SuSE, if someone wants to judge Linux as ready for the desktop, do not try Red Hat. Also, like SuSE, Red Hat works for experienced Linux users who want to spend a lot of money for a commercial distribution they have to tweak. My friends in that category laugh at me when they find out I use Red Hat. Now, I will laugh with them.

## GENTOO

The existence of Gentoo (<http://gentoo.org/>) became the reason for not including Slackware (<http://www.slackware.com/>) in my review. For all the Linux old-timers, Slackware meant Linux. Today, the Linux innovators use Gentoo.

For people just trying out Linux, keep away from Gentoo. As an US Army soldier once told me, they use Gentoo to hone their Linux skills. So, if you don't want to spend lots of time in boot camp, go a different direction.

For experienced Linux users, Gentoo has a wonderful desktop. Once you build your way to stage three, you might consider using emerge -k so you download binaries. However, if you want to compile everything, then just use emerge and wait a few days to get your system built.

For anyone wondering about what I'm writing, head over to the Gentoo web site and read the installation instructions.

## DEBIAN

Debian (<http://debian.org/>) represents many good things about Linux. People can obtain the distribution for free. It has many quality features. However, one needs to have Linux skills to install it and configure the distribution and the desktop.

Like Slackware and Gentoo, advanced users should consider Debian. If you want to create a high quality desktop, you'll need to do that yourself. While Gentoo provides a slick desktop, Debian relies on the user to engineer the desktop.

You might want to keep an eye on Debian. Bruce Perens, Director of the Linux Desktop Consortium (<http://www.desktoplinuxconsortium.org/>) has pointed to Debian as the new model of user-supported enterprise Linux. Bruce noted at the Desktop Linux Conference in Boston that a "user-driven Linux should involve multiple, independent support vendors and nothing held back from the free version of a distribution". He went on to say, "the function of distributions would be to add value as services for particular industries."

So, what might we see from Debian in the future? First, rumors that the team will port the Red Hat installer would put Debian in a new position among distributions. Secondly, a focus on hardware drivers and polishing consumer applications such as scanning and digital cameras could move to the forefront. Also, creating desktops for retail point of sales, ATM's, and thin clients would make Debian the leader in those fields.

Keep an eye on this distribution, you might wind up with it in a short time.

## XIMIAN DESKTOP 2 ON SuSE 8.2 AND RED HAT 9

If you do not have an inclination to go with Sun's desktop, Ximian's (<http://ximian.com/>) combined with SuSE will give you a decent desktop. Matthew Schick at the University of Southern Mississippi configures many Ximian Desktops on top of Red Hat and now SuSE. He gets excellent results.

Matthew maintains his own update servers. If you use

either SuSE's update service or Red Hat's up2date, expect glitches when you try to update. If you can live with some false dependency messages that don't have much effect on your desktop, you will find this an excellent solution.

With regard to electing packages, Ximian has done an excellent job of polishing OpenOffice 1.03. However, I did not select that option and downloaded Star Office 7 trial edition. I'm pleased that I did.

The StarOffice 7 trial (<http://www.sun.com/software/star/staroffice/get/index.html#eval>) provides a complete and functional suite. Sun did not disable anything. You can use it until it times out. I plan to buy it or purchase it with Sun's Linux Desktop.

I found it works well with Ximian 2. I also discovered that all the applications I want or need come free with Ximian's Red Carpet. I just subscribed to the Red Hat channel and check in with Ximian periodically. The updates will appear when you login to the network. Linux Ready for the Desktop?

Many people have argued back and forth about Linux for the desktop. As a person who spent a portion of his career grading the quality of "intelligence data", I want to share my perspective. While I do not practice law or work for the intelligence community, I still use those skills in my work as a technologist.

The information you read and people present about Linux, Windows and Mac lacks substance. The journalist who quotes some guy at IBM about the quality of the user interface has never programmed to the desktop libraries. I have programmed to those libraries.

Linux provides a superior set of desktop interfaces. Recently, Mac OS X finally added the UNIX Windowing to its Aqua interface. Apples says on their website (<http://www.apple.com/macosx/features/x11/>) that "prior to the introduction of Mac OS X, virtually the only way developers could create a graphical application in a UNIX-based operating system was with the X Window System, more commonly called X11. X11 for Mac OS X offers UNIX users the ability to run thousands of X11 applications concurrently with other Mac OS X applications". You might notice the "thousands of X11 applications" comment.

We also have Windows Services for UNIX (<http://www.microsoft.com/windows/sfu/default.asp>) version 3.0 which Microsoft says provides a full range of cross-platform services for integrating Windows into existing UNIX-based environments. Microsoft dedicates an entire homepage to this product offering.

I feel that most knowledgeable developers will say that Linux as well as other operating systems that use X11 as their graphical user interface are ready for the desktop. When I look at the Gnome 2.4 desktop available to Linux, Sun Solaris, Mac OS X and others, I see a mature, high quality and reliable desktop. The Real Difference

If IBM's Lotus division wanted to release all of its applications to Linux, it could in a heart beat. After all, many of those applications exist on UNIX X11 (<http://www.xfree86.org/>) desktops now. To get this done, IBM would have to settle the Internal divides within the company.

The hardware manufacturers could also release drivers and bundled applications for Linux too. Some have done so for MacOS X, why not Linux? One can only wonder since many hardware components only work on Windows and unlike the "I saved Christmas with my Mac" commercial, many drivers do not exist for the Mac.

So, applications which manufacturers refuse to port to Linux wind up being build by the Linux or, more accurately, the open source community. They have given us the [Apache web server](#), [OpenOffice](#), [Mozilla](#) and the [Netscape Browser](#), [the Evolution](#) personal management suite and [CodeWeaver's WINE](#).

I use them all and I also use CodeWeavers Crossover products which lets me run Win32 applications on Red Hat if I wish. So, if you want to run Photoshop on Linux, please go ahead. I have a couple of Win32 applications installed but don't use them except to show non-believers.

Also, one of the things you might find out has to do with Photoshop. On August 4, 2003, Anne Chen wrote in eWeek (<http://www.eweek.com/article2/0,3959,1210083,00.asp>): "Last year, when the Walt Disney Co.'s feature

animation unit, in Burbank, Calif., announced that it was using Linux for digital animation work, speculation grew that Adobe Systems Inc. would finally port its products to Linux. To this day, however, Adobe has done no such thing. Rather than wait, Disney, along with two other motion picture animation studios (which declined to be named for this article), decided to jointly fund the development of a Windows-to-Linux porting solution. The idea: develop technology using the Wine (<http://www.winehq.com/>) emulator to run Adobe Photoshop on Linux."

#### FINAL POINTS

I believe I titled this article appropriately. If someone evaluates the Linux desktop based on the applications it can run, then this product has everything it needs to succeed. If you evaluate the product on the applications it does run, you also have everything you need. If you run Windows 98 and have your eye on an upgrade path, consider Sun's Java Desktop made up of SuSE on the operating system side and the Gnome (<http://gnome.org/>) Desktop as your user interface. You'll save lots of money and discover happiness on your PC.

If you have some technical experience and feel like you can control your own environment, then look at the distributions mentioned above. If you feel like you need some help, contact Novell and tell them I sent you.

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

*<http://consultingtimes.com/desktoplinux.html>*



# AUUG Chapter Meetings and Contact Details

CITY	LOCATION	OTHER
<b>ADELAIDE</b>	TBA	Contact sa-exec@auug.org.au for further details.
<b>BRISBANE</b>	Inn on the Park 507 Coronation Drive Toowong	For further information, contact the QAUUG Executive Committee via email (qauug-exec@auug.org.au). The technologically 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.
<b>CANBERRA</b>	Australian National University	
<b>HOBART</b>	University of Tasmania	
<b>MELBOURNE</b>	Various. For updated information See:  <a href="http://www.vic.auug.org.au/">http://www.vic.auug.org.au/</a>	The meetings alternate between Technical presentations in the even numbered months and purely social occasions in the odd numbered months. Some attempt is made to fit other AUUG activities into the schedule with minimum disruption.
<b>PERTH</b>	The Victoria League 276 Onslow Road Shenton Park	
<b>SYDNEY</b>	Meetings start at 6:15 pm Sun Microsystems Ground Floor 33 Berry Street (cnr Pacific Hwy) North Sydney	The NSW Chapter of AUUG is now holding meetings once a quarter in North Sydney in rooms generously provided by Sun Microsystems. More information here: <a href="http://www.auug.org.au/nswauug/">http://www.auug.org.au/nswauug/</a>

FOR UP-TO-DATE DETAILS ON CHAPTERS AND MEETINGS, INCLUDING THOSE IN ALL OTHER AUSTRALIAN CITIES, PLEASE CHECK THE AUUG WEBSITE AT [HTTP://WWW.AUUG.ORG.AU](http://www.auug.org.au) OR CALL THE AUUG OFFICE ON **1-800-625655**.





**Application / Renewal  
Individual or Student Membership  
of AUUG Inc.**

Use this tax invoice to apply for, or renew, Individual or Student Membership of AUUG Inc. To apply online or for Institutional Membership please use <http://www.auug.org.au/info/>

**This form serves as Tax Invoice.**

Please complete and return to:

**AUUG Inc, PO Box 7071, BAULKHAM HILLS BC NSW 2153, AUSTRALIA**

If paying for your membership with a credit card, this form may be faxed to AUUG Inc. on +61 2 8824 9522.

Please do not send purchase orders.  
**Payment must accompany this form.**

**Overseas Applicants:**

- Please note that all amounts quoted are in Australian Dollars.
- Please send a bank draft drawn on an Australian bank, or credit card authorisation.
- There is a \$60.00 surcharge for International Air Mail
- If you have any queries, please call AUUG Inc on +61 2 8824 9511 or freephone 1800 625 655.

**Section A:**

**Personal Details**

Surname: .....

First Name: .....

Title: ..... Position: .....

Organisation: .....

Address: .....

Suburb: .....

State: ..... Postcode: .....

Country: ..... Phone Work: .....

Phone Private: ..... Facsimile: .....

E-mail: .....

Membership Number (if renewing): .....

**Student Member Certification**

For those applying for Student Membership, this section is required to be completed by a member of the academic staff.

I hereby certify that the applicant on this form is a full time student and that the following details are correct:

Name of Student: .....

Institution: .....

Student Number: .....

Signed: .....

Name: .....

Title: .....

Date Signed: .....

**Section B: Prices**

Please tick the box to apply for Membership. Please indicate if International Air Mail is required.

Renew/New* Individual Membership	\$125.00 (including \$11.36 GST)	<input type="checkbox"/>
Renew/New* Student Membership	\$30.00 (including \$2.73 GST)	<input type="checkbox"/>
Surcharge for International Air Mail	\$60.00	<input type="checkbox"/>

\* Delete as appropriate.

GST only applies to payments made from within Australia. Rates valid from 1st October 2002.

**Section C: Mailing Lists**

AUUG mailing lists are sometimes made available to vendors. Please indicate whether you wish your name to be included on these lists:

Yes  No

**Section D: Payment**

**Pay by cheque**

Cheques to be made payable to **AUUG Inc.** Payment in Australian Dollars only.

**OR Pay by credit card**

Please debit my credit card for A\$ .....

Bankcard  Mastercard  Visa

Card Number: ..... Expires: .....

Name on card: ..... Signature: .....

Date Signed: .....

**Section E: Agreement**

I agree that this membership will be subject to rules and bylaws of AUUG Inc as in force from time to time, and this membership will run from the time of joining/renewal until the end of the calendar or financial year as appropriate.

Signed: .....

Date Signed: .....

