

AUUGN

The Journal of AUUG Inc.

Volume 24 • Number 3
September 2003

Features:

Going 3D with Blender: A toy train	7
The Linux System Administrator's Security Guide	12
2003 and Beyond (part 2)	18
AUUGN Book Review: Google Hacks	29
Linux and Open Source in Government: CFP	30
AUUGN Book Review: Essential CVS	31
UNIX and Bioinformatics	32
Tuning and Optimizing Red Hat Advanced Server for Oracle9i Database	35
AUUGN Book Review: Drive-by Wi-Fi	49
Play Encoded DVDs in Xine	50
Interview: Australian Sun HPC Team	52
AUUGN 2003: The Photos	58


News:

Public Notices	5
AUUG: Corporate Members	31
AUUG: Chapter Meetings and Contact Details	63

Regulars:

President's Column	3
My Home Network	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

AUUG Management Committee

Email: 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.Lehey@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>

Malcolm Caldwell

Bus. Tel (08) 8946 6631, Fax (08) 8946 6630
<Malcolm.Caldwell@ntu.edu.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>

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>

Committee Members

David Bullock +61 402 901 228
<David.Bullock@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>

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>

Stephen Rothwell

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

Michael Still

12 Bushby Close Gordon ACT 2906
Phone: +61 414 382 568 <mikal@auug.org.au>

AUUG Business Manager

Elizabeth Carroll
AUUG Inc
PO Box 7071
Baulkham Hills BC NSW 2153
<busmgr@auug.org.au>

Editorial

Con Zymaris auugn@auug.org.au

A few words on the power of words. I am a great advocate of the need for technical people, AUUG people, to get out and speak their mind, to make their opinions, ideas and ideals heard by the great mainstream; by those whose utilisation patterns of technology we have foreseen and whose lives we are helping to shape the digital future of.

Words have the power to anger, the power to inspire. They have the innate attribute of providing a conduit for sharing ideas and conveying a vision which may not yet be in existence. Words, when sculpted by talented hands, are a joy to behold and contemplate. I came across one such pair of hands recently, belonging to one Ron Rosenbaum. Although not well known here in Australia, he has been called the quintessential magazine non-fiction journalist in his homeland. A native of New York, he has written for dozens of publications on a wide spectrum of topics.

One of his pieces appeared in a 1971 issue of *Esquire* magazine. Entitled "The Secrets of the Little Blue Box", it recounted Rosenbaum's travels through the eerily familiar geekscape of elite phone phreakers. Familiar, in the sense that almost every aspect of the mind-bending culture he found, interlinked globally by the public phone systems, eerily prophesied the global Internet culture which would follow two decades later. Everything from chat rooms, and security exploits to geeks like the fabled 'Captain Crunch', in search of enlightenment through their joining with the *system* as 'one'.

A few pearls spring to mind about this article. One is the use of the moniker *2600*, which has been adopted by hacker culture ever since. At some point, it was realised that AT&T's (*Ma Bell* to the inhabitants of this realm) technology had an unlocking mechanism which made available access to the system functions of the phone system, when a tone of 2,600 cycles per second was played into the mouthpiece of a phone. It just so happens that this exact tone was generated by whistles thoughtfully provided by the cereal manufacturer which produced the Cap'n Crunch breakfast cereal, thus the appellation that John Draper donned to become a proto-hacker legend.


Now, what was so important about Rosenbaum's words that it's worth focussing attention on these handful of pages written over three decades ago? Indeed, why should we, as technical people, take the time and effort to cement our ideas in writing and invite others to read and contemplate them? When the piece in *Esquire* was published in 1971, a couple of guys in far-off California read it. They were blown away by what the geeks highlighted in the article had done, and decided that they too wanted to hack phone systems and electronics. They set about trying to manufacture these illicit devices in their parents' garages. Who were they and why did it matter? Oh, just a couple of guys called Steve... Cheers, Con

Contribution Deadlines for AUUGN in 2003

Volume 24 • Number 4 – December 2003: **November 15th, 2003**

Volume 25 • Number 1 – March 2004: **February 15th, 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

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>, Rik Harris <rik@kawaja.net>, Raymond Smith <zzrasmit@uqconnect.net>, David Lloyd <lloy0076@adam.com.au>, Peter Sandilands <peter@sandilands.vu>, Grahame Bowland <grahame@ucs.uwa.edu.au>, Cameron Strom <c.strom@statscout.com>, Steve Jenkin <sjenkin@canb.auug.org.au>

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>

DEATH OF THE INTERNET, FILM AT 10

For many of us, the Internet has been part of our lives for years, even decades. During that time, it has made the transition from a tool for privileged researchers to a playground for the general public, including kinds of people many of us never knew existed.

The change has been interesting, and it hasn't been without problems. The lawmakers initially ignored it, cheered on by the old-timers in the assumption that the lawmakers, knowing nothing about technical matters, would legislate badly.

Looking back over the last couple of years, things have changed. Yes, we have had our legislation, some of it confirming our assumptions about the lawmakers, but we have also had too many non-lawmakers who have either misunderstood the medium or sought to exploit its lawlessness. There seem to be a number of categories. The users include:

"Us", both the old-timers who were using the Internet before it became a public buzzword, and also the technically literate newer users who understand the difference between domain names and "web addresses".

- ◆ The users. Most computer users are functionally illiterate. For them, using computers is a matter of knowing which buttons to press (sorry, "click"). It's not based on understanding.
- ◆ A lot of people are in the Internet industry for business, of course. That includes "us" again, but also:
- ◆ The cowboys, people who consider themselves (relatively) honest, but who are out to make a quick buck off the Internet. For some reason, domain registrars seem to show a remarkable proportion of such people.
- ◆ The criminals, people who are out to make a quick buck off the Internet, cost what it may.
- ◆ The software vendors who exacerbate the problems with buggy software. I don't need to name the prime suspect.

We see the results of the problem every day: NOIE, the National Office of the Information Economy, estimates that half of all Internet traffic is spam. It would be easy to surmise that the rest is probably "viruses" and attempts to break into other systems. I have discovered that I can halve my Internet traffic, without missing anything I want, by the simple expedient of reducing the number of IP addresses routed to the Internet from 256 to 16: the 87.5% reduction in unsolicited traffic attempting to break into the network makes more than half my total traffic.

Why should this be allowed to continue? Nobody

wants it to, of course, but everybody points at somebody else. The uneducated users blame Microsoft. Educated users fix their installations, either by installing Microsoft security updates (not the spurious ones circulating as part of yet another exploit), or better, by installing a real operating system.

Microsoft points to the difficulty of fixing security problems, since there will always be new ones. It's amusing to compare the resources available to them with those available to the OpenBSD project, arguably one of the most secure operating systems available. The OpenBSD project has very little money and only a little over 100 developers, nearly all unpaid, but they produce a system incomparably more reliable than Microsoft's offerings. It's difficult to take Microsoft's claims seriously.

One problem remains, of course: Microsoft can issue as many security fixes as it want, but they won't have any effect until they're installed. Many casual Microsoft users, myself included, can't be bothered to install them. People like this should not be allowed to connect to the Internet (and yes, I firewall off my Microsoft addresses to ensure that this can't happen).

But how do you stop people from abusing the Internet, either deliberately or accidentally? With very few exceptions, the people who cause the problem won't take any action, either out of ignorance or out of malice. The only people who can ensure that they don't get any access are the ISPs.

The ISPs, of course, point at the impossibility of the task. It would take an incredible effort to police their users. They might need a multiple of their present staff just for this function alone. How can they do that?

The objection seems reasonable. I'll get back to it in a while. But first let's take another look at what has been happening lately. In August, a large part of the North-East of the United States suffered a power blackout, the largest that country (and thus presumably the world) has ever seen. I've seen a couple of news reports blaming the blackout on security breaches in Microsoft software. The Microsoft bashers amongst us had fun with that one, of course.

The report seems to be incorrect. Does that make any difference? There are two reasons to think it doesn't: firstly, in today's networked environment, exploits of this nature are both becoming more common, and secondly their effects are becoming more significant. Every new Microsoft exploit costs millions of dollars. Many experts consider that the current situation is just a catastrophe waiting to happen.

So, again, why isn't anything being done? It is. The lawmakers are finally becoming active: in September, Senator Alston introduced the "Spam Bill 2003", which will make a number of current practices illegal. That's a step in the right direction,

but it's not enough. One of claimed reasons for lack of further action is the anonymity of the Internet. That's not enough: a similar problem exists with tax evasion. Tax offices world-wide expend significant effort chasing tax evaders, and they're largely successful. Given the log records associated with spam and viruses, it should be easier to address them. The only conclusion that I can draw is that people don't want to address the issue.

That brings us back to the poor ISPs. Their objection is valid: policing access in this manner would take a lot of work. That's not a reason not to do it, though. If I drive my car into a shop window, I can't absolve myself of blame because the brakes weren't working. By contrast, people with broken computers can connect to the Internet and then blame their computers for the damage they do.

For ISPs to be able to afford to police connections, things need to change. Somebody has to pay for the work they do. Currently, nobody will: if an ISP doubles his prices to pay for the work, he will lose customers and go broke. It's equally impractical for an individual ISP to introduce an AUP which imposes draconian fines on people who, possibly accidentally, spread spam or viruses into the Net. This, too, would probably cause the ISP to go broke.

I can only see two alternatives here: either the ISP industry as a whole (world-wide!) must come to a code of practice which includes such policing, or it must be legislated. The former would work better: I'd guess that if the ISPs responsible for 90% of the world's Email were to come up with an enforceable policy of disconnecting spammers and virus relays from the net, they would be in a position to force the remaining 10% to either adopt the policy or be blocked from sending mail to them. Will it happen? I'm sceptical.

The alternative is legislation. If there's one positive thing that has come out of the mess of the past two years, it's the realization that Something Must Be Done. Politicians are now also being harassed by the situation. They have the double problems of limited

understanding and limited geographical scope, but it would be a start.

silicon breeze



Public Notices

Upcoming Conferences & Events

LISA

5th Symposium on Operating Systems Design and Implementation
December 9-11
Boston, MA

Linux and Open Source in Government

The challenges
12-13 January 2004
Adelaide,

linux.conf.au

Jan 14 until Jan 17
Adelaide

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 improving 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 3rd parties (Cybersource, or others) to undertake this work for you.

The fee schedule for allocPSA includes a pre-purchased quantity 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 

My Home Network (September 2003)

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

Well, my first column in a few months, and well after all the regular activities such as conference, elections, etc. All up it is the end of a very busy time. In recent columns, I've been writing about low level network setups, and other system level items, but this time I'm going to tackle a very different end, the user interface (UI) and commonly used tools and facilities.

As many people know, for some time, I've been using Red Hat 9 (or earlier versions), but often for simple connections have used that dreaded Microsoft product, such as Outlook Express and Internet Explorer. The main reason for this is that the most common attachment to mail is a Word Document. Until recently, such tools as Abiword, while adequate were not the greatest translation. With the development of OpenOffice this has change dramatically.

OpenOffice provides not only a word processor which can handle Microsoft Word documents, but also Excel spreadsheets and Powerpoint presentations. In addition, it is not only available for Linux, but also most Unix versions and even MacOS X and the various Microsoft Windows systems.

Rewinding a bit and explaining about my choices in UI. For many, many years I believed that a simple text interface was fine and such tools as `mailx` and `vi` were fine for most uses. In fact, I started on `ed`, so even a curses interface is not essential. However, over the years people have started sending various attachments to their mail, so `mailx` had to be replaced by `mutt`, which, while still a text based mail reader, using curses, is able to handle attachments. If I really wanted to be picky, `lynx` and similar interfaces are often adequate for many web pages.

Unfortunately, the computing world keeps moving on, or at least changing anything that works, so there are times when I need to use a Graphical User Interface (GUI). What is worse, GUI's are addictive, the more you use it, the more things you find that do use it.

Fairly obviously, the basis of any GUI on Unix and Linux systems is X11. and just as obviously, on Linux it is XFree86 based. X11 is only the lowest level, with the biggest battle for the GUI being in terms of the Window Manager. Over the years there have been many different ones, but today it has generally come down to a "battle" between GNOME and KDE. Unfortunately, the choice between the two is based more on religion than rational decision, as most programs can be run under either.

My preference is for GNOME, mainly because it has a simpler "look and feel", and more importantly, it is standard under Red Hat. In fact, my personal opinion is that it is probably better to stick with your installation defaults, rather than make a religious

decision. For example if you like KDE, but want a Red Hat style distribution, then an obvious choice is Mandrake.

The GUI itself is just an enabler, and the more important thing is what you do with it. Well, the first thing that forced me to move away from a text interface was a web browser. As the most common application used by the world for surfing the "web", everyone expects you to view all sorts of pictures and other non-text items. To handle it, very early on, I began using Netscape on various Unix platforms, but these days use Mozilla. This isn't much of a stretch, but while there are many competitors to Mozilla, it is simple and easy to use.

By itself, Mozilla is a fairly powerful web browser, but one of the advantages of any browser these days is the use of plug-ins, which extend the features available. As an average web user, the requirements are not too extensive, but some extensions are needed.

The most useful one I have is 'plugger', which acts as a streaming multimedia plug in for Mozilla, but more importantly, allows normal Unix programs to be run against downloaded data. For example, MPEG files will be passed to a series of programs, including, 'xine', 'mplayer', 'mtvp' and lastly 'xanim', with different arguments. If any of the programs aren't found, then it will look for the next.

Of course, not everything can be handled in this way. The most common plug-in, particularly in the Microsoft world, is Macromedia Shockwave Flash. This is also available for Mozilla, and is a simple download from the Macromedia site (www.macromedia.com). While it can sometimes be a bit behind the one for the Win32 world, it is generally able to handle most of the latest Flash downloads.

Of course, no browser is complete these days, without Java, and since Sun is primarily a Unix company, Java plug-ins for Linux are easy to come by. They can be installed either for an individual or system wide. For system wide use, you just need to install the file 'libjavaplugin_oji.so' into '/usr/lib/mozilla/plugins/'. It will then invoke Java and pass any Java calls directly to the Java VM and return the results.

The final plug-in I have installed which isn't a major item is VLC is again mainly used to handle various video and audio applications. It is a by product of the installation of VLC for playing DVD's.

While I am using Mozilla as a web browser, I don't use it for a mail reader. It certainly is fully functional, but I've found separating the two is better. For my environment, I much prefer to use Evolution, which originally was with Ximian distributions, and is a groupware suite rather than just mail reader. Basically, it started out to be a Outlook Express "look-a-like" for Linux (there is even a connector for Microsoft Exchange available), but now seems to have taken on a life of its own.

Aside from handling all standard mail access protocols (i.e. Mbox, MailDir, POP and IMAP) it also acts as a calendaring program, address book and a task-list. The address book can connect to an LDAP server. Like most such applications, it can be used to connect to multiple different servers and accounts, and this proves useful as I'll show below.

For reading news, i.e. NNTP, there are a number of options, but the one I've found simplest is 'pan'. It is very simple to set up and use. However, useful Usenet news streams are becoming rarer and rarer, but there are a couple of ones I do enjoy, such as 'rec.humor.funny'.

One of the latest additions to my list is 'gaim', one of the most active projects on SourceForge, and an excellent Instant Messaging client. Unlike Microsoft products like MSN Messenger, 'gaim' allows access to a number of different servers, including AOL's Instant Messenger (AIM), ICQ, IRC and many others. Many of these are implemented as a plug-in to 'gaim' and not all are as flashy as the Win32 equivalents, but has similar functionality. In the past I used 'everybuddy', but have since switch to 'gaim' due to many of its features.

More recently, I've added a DVD reader (and writer) to my system and would like to watch DVD movies. While this is a complex legal area, for Linux uses it isn't such a big issue. There is only one product that really suits, that is VideoLAN Client, from <http://www.videolan.org/vlc>. Like Mozilla, it is available for many different platforms, including Microsoft Windows, BeOS and Solaris. In fact, it is without doubt the best freely available DVD player on any OS.

The one issue I had with VLC was that there are a large number of other packages that need to be installed along with it. While a pain, it is reasonable, as there are many different video and audio codex that need to be supported. It is also based on a number of other low level libraries and other DVD support packages (e.g. libdvdplay, libdvdcss, libdvdread, etc) needed to handle access to DVD structures. While there are a number of other packages, these all have links from the VLC site. Once installed, just loading a DVD will cause autoplay to run 'vlc' and play the DVD movie.

As I started describing, I'm generally happy with a text based interface, and so one of the most heavily used programs I have is 'gnome-terminal'. In fact, in recent times, I've found a few new features, such as 'tabbed' terminals, i.e. where a number of different terminal sessions occupy the one terminal window. Very useful when you have related items, but don't need to always deal with one of the terminals.

The final utility I find useful is ssh, or more specifically OpenSSH, which allows secure connection to various sites. More importantly, it can also be used to forward TCP ports across that secure tunnel. This can be used to securely connect to remote IMAP

servers, web servers and other connection end points. For example, I have the following in the file `~/ .ssh/config` on my home machines:

```
Host *
  StrictHostKeyChecking ask

Host sc.apac.edu.au
  User fpc561

Host www.auug.org.au
  ForwardX11 no

Host frank*
  LocalForward 1430 203.202.1.130:143
  LocalForward 3128 203.202.0.69:3128
  LocalForward 5900 192.168.3.4:5900
  LocalForward 5901 192.168.3.4:5900
  LocalForward 3389 127.0.0.1:3389
```

Put simply, this connects to an IMAP server, a Squid cache server, two different VNC servers and a rdp (i.e. Microsoft Windows Remote Desktop) server on any host named `'frank'`. In addition, for one `sc.apac.edu.au`, I change my username and for `www.auug.org.au` I don't support X11.

One interesting point to note is that, as a normal user you cannot map privileged ports, i.e. anything < 1024 (e.g. IMAP's 143), you can do remapping of a non-privileged port to a privileged port, as specified for IMAP. I'll leave it as an exercise for the reader to add a line to forward outgoing connections back the required SMTP server.

To use these mapped ports, it is necessary to connect to the ports on the localhost. For example, for Evolution, the destination host is specified as `"localhost:1430"`. For VNC, the two ports are accesses as `"localhost:0"` and `"localhost:1"` (i.e. VNC adds 5900 to the specified port number).

Finally, to connect to the rdp host I need to on-forward port 3389 on the remote host onto another host. However, it is included here more as a reminder rather than a used redirection. Once `'ssh'` is running, additional ports can be added with the escape sequence `~C` to add additional "command-line arguments", e.g. use the sequence:

```
~C:
~L:3390:203.204.1.130:3389
```


to add an additional port redirection.

As an additional benefit of using these applications on Unix/Linux, it is easy to keep different systems in sync, by just copying around text files from the appropriate directories rather than trying to dump and restore data from a binary registry.

Of course, all that I've listed is just my choices of what to use on my desktop, you may well use something else, or make some other choice in terms of Window Manager, browser or any other application. If so, I'd like to hear about it, and so would others.

Going 3D with Blender: A toy train

Author: Katja Socher <katja@linuxfocus.org>



ABSTRACT

In this second article in our series about modelling with Blender we create a little toy train. We assume that you have read the first article *Going 3D with Blender: Very first steps* and built that little stage yourself as we are now using it as our starting point.

GOING 3D WITH BLENDER: A TOY TRAIN

Look at the illustration picture above and imagine it without the Tux penguin. We are now going to create a similar toy train.

Open the default.blend file with our stage settings that we created last time as the starting point for our toy train. (In case you saved them with the sphere and the cube select them by a right click while holding shift down and press `x` to delete them.) We will begin with modeling the last wagon. I still used Blender version 2.27 when writing this article.

BUILDING THE LAST WAGON FIRST

To build it we start with a cube (for the green part): In front view press `Space` and then `Add`→`Mesh`→`Cube`. Then press `tab` to leave edit mode. With the cube still selected press `s` (in front view) and scale the cube down so that it has the right height. For the length press `s` again and hold the middle mouse button down while moving the mouse to the right. This way you restrain the change in size to the side you are moving the mouse. Finally for the width go to side view, press `s` and again move the mouse to the right while holding the middle mouse button down.

Sometimes you will find that Blender doesn't want to do what you intended to do. So you also have the possibility to work with numerical values. Just hit `n` to get to the menu and change the `x`, `y` and `z` values accordingly.

To give it a green colour go to the material button and the white button, click "add new" and change the colour to green (to get exactly my colour: `R=0, G=0.82`

and B=0). You can either move the sliders with the mouse or you can left click on the letters and then type in the values.

Now we need wheels for our wagon. In front view press Space, then Add→Mesh→Cylinder (leave the vertices at their default value of 32) and tab to leave edit mode. Make the cylinder a bit smaller and thinner as you did with the cube: in side view first press *s* and scale the whole cylinder down then press *s* again and hold the middle mouse button down while moving the mouse to the right (or hit *n* to use the numerical values) to make the cylinder thinner. Give it a red material (go to the material button, then press the white button and "Add new" and move RGB to 1, 0 and 0 respectively), then in top and front view place it (press *g* and move the mouse) on one side of the wagon. The first wheel is ready.

Copy it by pressing shift + *d* and move (press *g*) the second wheel to its place. Repeat this until you have all four wheels in place.

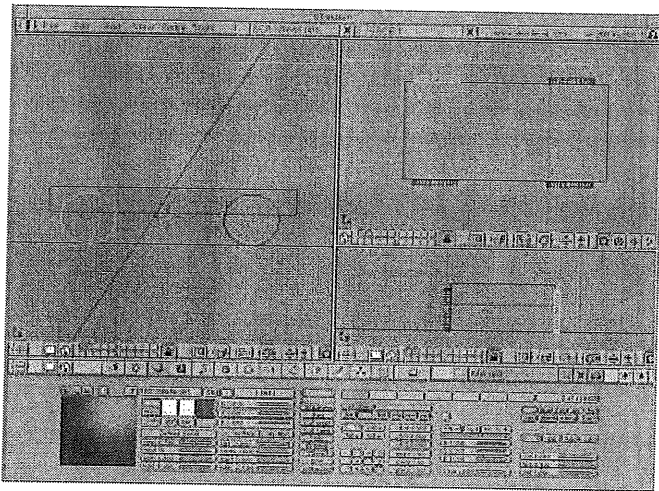


Figure 1: Wheels and green part of the last wagon

Now we still need to build the yellow load of the wagon. In side view hit Space then Add→Mesh→Cylinder, then tab and in top view scale it down (press *s*) so that the width is as big as the width of the green part. Now press *s* again and restrain the scaling to the length again (move the mouse to the right while holding the middle mouse button down or hit *n* to use numerical values). The load should fit exactly on the wagon. Now press *g* and in front view place the load on the wagon if you haven't done it already and by going to the material button, clicking "Add new" and changing the colour sliders to R=1, G=1 and B=0 you change the colour of the load to yellow. The first wagon is ready! Congratulations!

THE DARK BLUE WAGON

Next is the dark blue wagon which is quite easy to build as you probably already see. In front and top view just select the green cube and the four red wheels by right clicks while holding shift down (in case some other part is also selected unselect it by another right

click while still holding shift down) and then duplicate it by pressing shift + *d*. In top view move it to the left and place it besides the other wagon. In front view change the height of the wagon by right clicking on the cube, then pressing *s* and moving the mouse upwards a bit while holding the middle mouse button pressed. Now move the whole cube a bit up and give it a dark blue colour by going to the material button, click "add new" and move the RGB sliders to blue (R=0, G=0, B=1).

It might be a good idea to name the materials according to their colour. So left click in the field that starts with MA: and replace the default Material.00x by typing blue. Do the same for the red, yellow and green materials (e.g. select a wheel and go to the material button, the red colour is shown now, left click in the field that starts with MA: and replace the default Material.00x by typing red and so on.)




Figure 2: Naming the blue material "blue"

By the way you can use + and - of the numpad to zoom in and out of your views. If you want to change the section that is visible after you have zoomed in for example you can move within the view by pressing shift and the middle mouse button while moving the mouse.

Depending on how big you have made the two wagons it can be necessary to scale them down when you add more wagons as you won't be able to see them on your rendered image otherwise. Just select all objects that belong to the train either by pressing *b* and marking (draw a rectangle around) the two wagons (shift + right click on objects you just unintentionally selected by this) or by holding shift down and clicking on every little object individually, then press *s* and change the size of all wagons. This way all wagons will get the same change in size and will so fit nicely to each other.

The wheels of our wagons should also slightly touch our floor. By scaling them down they have probably moved themselves up a bit so press *g* and move everything down on the floor again.

You can always press F12 to make a render and see if the train looks good.

THE ORANGE WAGON WITH RED AND GREEN CYLINDERS

The third wagon is easy as well: Again select the green wagon and its four wheels (press *b* and mark the objects (draw a rectangle around them), make sure that only the objects you wanted to select are really selected, otherwise unselect them by shift + right click), duplicate it (shift + *d*) and move it to the left of the dark blue wagon. Change the colour of the wagon to orange (right click on the cube to select it, then go to the material button, press "Add new" and change the

colour sliders to R=1, G=0.647, B=0, don't forget to name your material "orange" then). Now in top view click with the red-white cursor in the middle of the right part of the wagon, then press Space then Add→Mesh→Cylinder, then tab to leave edit mode. Scale the cylinder down (press s) and then in front or side view place (press g) the cylinder on top of the wagon and scale the height up (press s again with holding the middle mouse button down to constrain the scaling) and give it a red colour (go to the material button and click the white button and then red).

The cylinder should now be in the middle of the right part of the wagon (you can see this e.g. in top and front view). Next you only need to duplicate it (shift + d), place it (press g) in the middle of the left part of the wagon and change the colour to green (go to the material button and choose "green" from the list of materials).

Now the third wagon is ready too!

THE RED AND GREEN WAGON

For the fourth one select all parts of the dark blue wagon, duplicate them (shift + d) and move them to the left. Select the dark blue cube (right click) and change the colour to green. Now in front view duplicate it and place it on top of the other. Change the colour to red. That's it already.

THE ORANGE WAGON

The same for the next one: In top view select all parts of the dark blue wagon (right click on the blue wagon, then press b and mark the objects), duplicate them and move them to the left. Now select the dark blue cube (right click) in front view again and change its height and its colour (to orange). As the cube will move a bit downwards by scaling it up, move it a bit up. Here we go.

THE LIGHT BLUE WAGON

By now you should have some practice with moving, scaling, rotating and changing colours. So for the last wagon we will use a new technique.

As you can see it is different from the others as it is open on top. To create it add a grid in top view (hit Space, then Add→Mesh→Grid). You are asked for Xres and Yres, choose 8 with both. Press tab twice and select the two outer lines of all four sides (press b and mark the two lines on top, then press b again to mark the next two lines until you have selected all the two outer lines of the grid).

Figure 3: The grid with the two outer lines of all four sides selected

Now in front view extrude the selected points by pressing e and moving the mouse upwards. By extruding you create three-dimensional geometry out of flat, two-dimensional shapes and it is a technique you will frequently use when you work in 3D.

The rest you should already know by now: press tab to leave edit mode. Give it the right size (press s and scale it to make it fit to the other wagons) and a light blue colour (R=0 G=0.714, B=1), add the four wheels by copying them (shift + d) from another wagon and place them accordingly (press g).

Finally you hit Space and then Add→Mesh→UVsphere (change both, the number of the segments and the rings back to 32) to add the sphere. Press tab, scale the sphere down (press s), place the sphere inside the wagon (press g) and change the colour to pink (go to the material button, then click the white button and "Add new", then move the colour sliders to R=0.8, G=0 and B=1). Now go to the edit button (the button that looks like a square with yellow edges) and press "set smooth" (as you already did with the sphere in our previous article). The sphere will get a much smoother look (a smooth surface).

THE CONNECTION BETWEEN THE WAGONS

Before modeling the locomotive let's first connect our wagons: Again add a grid (I chose 8 for the values of Xres and Yres again but you could also choose Xres=8 and Yres=2) in top view to our scene, extrude it a little bit in front or side view by pressing e and moving the mouse, then press tab to leave edit mode and now in top view again resize it (press s). Its length should be from the center of one wagon to the center of the other with a little space between the two wagons and it shouldn't be larger than one-third of the width of the wagons (see Figures 4 and 5). Give it a material (R=0, G=1, B=1). Duplicate it (shift + d) and place it between all wagons (don't forget the one between the last wagon we modeled and the locomotive).

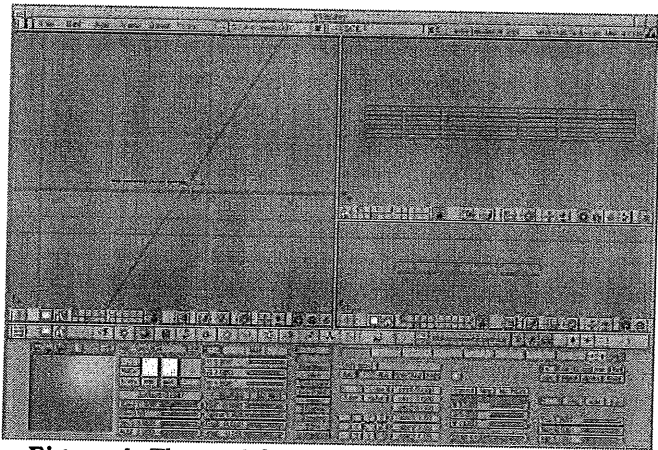


Figure 4: The grid for the connection between the wagons





Figure 6: The outline of our locomotive




Figure 5: A connection grid between a wagon

THE LOCOMOTIVE

Now the locomotive shouldn't be too difficult either: Again add a grid (Xres and Yres being 8 again) to our scene in top view. Then extrude it (press e) in front view. Press tab twice and select (press b, then mark with the mouse) only the right upper half of the locomotive (the four points on top and to the right). Extrude (press e) this part upwards again. Now give the whole object a yellow colour.

The four upper points should still be selected, extrude (press e) them again (but only a little bit). The new upper points are selected now. Still having them selected press s while holding shift down and move the mouse away from the model to make this part bigger (see Figure 6). Now we have the outline.

To give the roof of the locomotive the blue colour select the first line of points below them as well (at the moment only the top upper line is selected, to select the line below them as well press b and mark the line as before). Then go to the edit buttons and press "new" under the box where the colour is displayed. Next press the "Select" box. Before pressing "Assign" go back to the material button again and click on the blue material. (Note that if you wanted to change the colour to one that doesn't already exist you first have to press "Add New" before you change the colour sliders to your chosen colour.)




Figure 7: The buttons to assign more than one colour to one object

The rest is "old stuff" again: Scale the locomotive to a size that it fits to the other wagons (press s), add the wheels like for all the other wagons and place the locomotive to the left of the other wagons. Then duplicate (shift + d) the red cylinder from the wagon with the two cylinders, resize it (press s) and place it in front of the locomotive. Duplicate and resize this again (it should have half the length of the other), move it to its place and you have successfully modeled your toy train!

Here is a screen shot of the three different views in Blender:

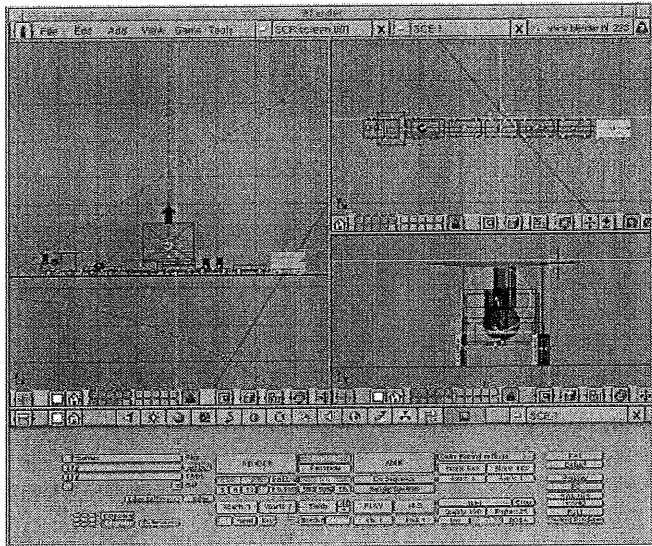


Figure 8: Our train in the three different views

FINALLY

Your toy train is ready now. But somehow if you render it (press F12) no shadows are displayed (even though you have pressed the Shadow button in the Display menu). Just select the spot light and press the "Only Shadow" button in the lamp button menu. Now give the spot light a higher energy level (e.g. a value of 5.0) and render your image again. There should be some shadows visible now.

Time to admire your work! :)

Here is our train:




Figure 9: Our train

If you move the camera a bit up in side view and then rotate it so that it looks down on the train you get this:




Figure 10: Our train from above

To save it as a .jpg image enter the directory and file name in the field with the default entry /render (left click, then type in the name) and also press the Extensions button. With this button pressed the picture is actually saved with the extension .jpg (or whatever format you choose) while otherwise the extension is not shown. Press the OSA button plus any of the numbers below (for quality), press the shadows button to see the shadow the train creates, change the End:250 field to End:1 (either make a left click and move the mouse to the left until the number has changed to 1 or hold shift down while left clicking in the field and then type in the value), choose the values of SizeX and SizeY depending on how big you want your image to be, choose Jpeg or any other format, press the RGB button and when you now hit the anim button your rendered image is finally saved.

Don't forget to save your train as a .blend file as well (go to the menu→Save as→type toytrain.blend (or however you want to name it)→Save file) so that we can reuse it again next time!

Have fun and happy blending:)

REFERENCES

- The Official Blender site (here you get the latest information about the further development of Blender, you can download it, there are tutorials): <http://www.blender.org>
- Blender cafe (in English and French): <http://www.linuxgraphic.org/section3d/blender/pages/index-ang.html>
- General articles about 3D graphics and animation: <http://webreference.com/3d/>

Katja is the German editor of LinuxFocus. She likes Tux, computer graphics, film & photography and the sea. Her homepage can be found <http://www.toppoint.de/~utuxfan/kj/>.

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

<http://www.linuxfocus.org/English/September2003/article307.shtml>

AUUGN CD

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

Due to circumstances beyond our control, this edition of AUUGN does not include a CD-R. We apologise for the omission. To make up for it, the next edition will contain two CD-Rs. As always, we're interested in your suggestions for CDs to put in AUUGN. Contact Liz Carroll at busmgr@auug.org.au if you have an idea.

Linux.Conf.Au 2004

Author: Michael Davies <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/>

The Linux System Administrator's Security Guide

Author: Kurt Seifried <kurt@seifried.org>

EDITORS NOTES AND LICENSE

This is a serialization of Kurt Seifried's Linux System Administrator's Security Guide. Each AUUGN edition will contain two to three sections (depending on space) from the guide. This is the second part of the installmant.

ADMINISTRATIVE TOOLS

There are a variety of tools to make administration of systems easier, from local tools like sudo which grant limited superuser privileges to www based systems that allow for remote management from a cybercafe while on vacation. For information on how to login remotely (i.e. interactive shell prompts) please see the [shell server](#) section.

LOCAL TOOLS

While it is possible to administer a Linux system from the command line using no "additional" tools it can be bothersome. If you wish to split up administrative tasks the "sub administrators" will often require root access to restart daemons, modify configuration files and so forth. Simply giving them all root access, or sharing the root password is often the first step to serious problem (this is one of the major reasons many large sites get broken into).

YAST

YaST (Yet Another Setup Tool) is a rather nice command line graphical interface (very similar to scoadmin) that provides an easy interface to most administrative tasks. It does not however have any provisions for giving users limited access, so it is really only useful for cutting down on errors, and allowing new users to administer their systems. Another problem is unlike Linuxconf it is not network aware, meaning you must log into each system you want to manipulate. YaST version two is now available and includes many new features as well as bug fixes, it is recommended you upgrade.

SUDO

Sudo gives a user setuid access to a program(s), and you can specify which host(s) they are allowed to login from (or not) and have sudo access (thus if someone breaks into an account, but you have it locked down damage is minimized). You can specify what user a command will run as, giving you a relatively fine degree of control. If you must grant users access, be sure to specify the hosts they are allowed to log in from when using sudo, as well give the full pathnames to binaries, it can save you significant grief in the long run (i.e. if I give a user sudo access to "adduser", there is nothing to stop them editing their path statement, and copying bash to /tmp/adduser and grabbing control of the box.). This tool is very similar to super but with slightly less fine grained control. Sudo is available for most distributions as a core package or a contributed package. Sudo is available from <http://www.courtesan.com/sudo/> (just in case your distribution does not ship with it). Sudo allows you to define groups of hosts, groups of commands, and groups of users, making long term administration simpler. Several /etc/sudoers examples:

```
#Give the user seifried full access
seifried ALL=(ALL) ALL
#Create a group of users, a group of hosts, and
allow them to shutdown the server as root
Host Alias WORKSTATIONS=localhost, station1,
station2
User Alias SHUTDOWNUSERS=bob, mary, jane
Cmdnd Alias REBOOT=halt, reboot, sync
Runas Alias REBOOTUSER=admin
SHUTDOWNUSERS WORKSTATIONS=(REBOOTUSER) REBOOT
```

SUPER

Super is one of the very few tools that can actually be used to give certain users (and groups) varied levels of access to system administration. In addition to this

you can specify times and allow access to scripts, giving setuid access to even ordinary commands could have unexpected consequences (any editor, any file manipulation tools like chown, chmod, even tools like lp could compromise parts of the system). Debian ships with super, and there are rpm's available in the contrib directory. This is a very powerful tool (it puts sudo to shame in some ways), but requires a significant amount of effort to implement properly (like any powerful tool), and I think it is worth the effort. Some example config files are usually in the /usr/doc/super-xxxx/ directory. Super is available <ftp://ftp.ucolick.org/pub/users/will/>.

WWW BASED TOOLS

WWW based administration tools provide an attractive solution since virtually every modern computer and Internet access point is web capable (sometimes that is all they are capable of).

WEBMIN

Webmin has had number of security problems so make sure you are using the most recent one. Webmin is one of the better remote administration tools for Linux, written primarily in Perl it is easy to use and easy to setup. You can assign different 'users' (usernames and passwords are held internally by Webmin) varying levels of access, for example you could assign bob access to shutdown the server only, and give john access to create/delete and manipulate users only. In addition to this it works on most Linux platforms and a variety of other UNIX platforms. The main 'problem' with Webmin is somewhat poor documentation in some areas of usage, and the fact that the username/password pair are sent in clear text over the network (this is minimized slightly by the ability to grant access to only certain hosts(s) and networks). Most importantly it makes the system more accessible to non-technical people who must administer systems in such a way that you do not have to grant them actual accounts on the server. Webmin is available <http://www.webmin.com/webmin/>, and is currently free. Webmin defaults to running on port 10000 and should be firewalled.

LINUXCONF

Linuxconf is a general purpose Linux administration tool that is usable from the command line, from within X, or via it's built in www server. From within X it provides an overall view of everything that can be configured (PPP, users, disks, etc.). To use it via a www browser you must first run Linuxconf on the machine and add the host(s) or network(s) you want to allow to connect (Conf > Misc > Linuxconf network access), save changes and quit. Then when you connect to the machine (by default Linuxconf runs on port 98) you must enter a username and password. By default Linuxconf only accepts root as the account, and Linuxconf doesn't support any encryption (it runs standalone on port 901), so I would have to recommend very strongly against using this feature across networks unless you have IPsec or some other form of IP level security. Linuxconf ships with several distributions and is available

<http://www.solucorp.qc.ca/linuxconf/>. Linuxconf also doesn't seem to ship with any man pages/etc, the help is contained internally which is slightly irritating.

OTHER NETWORK BASED TOOLS

On the other hand web based administration tools tend to be limited, and are typically not designed for heterogeneous installations (i.e. Linux, HP-UX, AIX and so forth). "Industrial" strength tools may be called for, like the following ones.

PIKT

Pikt is an extremely interesting tool, it is actually more of a scripting language aimed at system administration than a simple program. Pikt allows you to do things such as killing off idle user processes, enforcing mail quotas, monitor the system for suspicious usage patterns (off hours, etc), and much more. About the only problem with Pikt will be a steep learning curve, as it uses its own scripting language, but ultimately I think mastering this language will pay off if you have many systems to administer (especially since Pikt runs on Solaris, Linux and FreeBSD currently). Pikt is available at: <http://pikt.uchicago.edu/pikt/>.

VNC

Virtual Network Computer (VNC) is similar to X or PCAnywhere. You can display a graphical desktop, and control it remotely, with NT or Linux as the server and/or client. VNC across 10 megabit Ethernet is quite good, however it does tend to use a lot of computer power relative to other methods of remote administration. You can get VNC <http://www.uk.research.att.com/vnc/>. Security VNC isn't so great, but there are several sites with information on securing VNC, using SSL, SSH and other methods. There is also a page on securing VNC with SSH port forwarding at: <http://www.zip.com.au/~cs/answers/vnc-thru-firewall-via-ssh.txt>.

CFENGINE

cfengine is a set of tools for automating administration tasks and is network aware. You can get cfengine <http://www.cfengine.org/>.

BACKUPS

I don't know how many times I can tell people, but it never ceases to amaze me how often people are surprised by the fact that if they do not backup their data it will be gone, if the drive suffers a head crash on them or they hit 'delete' without thinking. Always backup your system, even if it's just the config files, you'll save yourself time and money in the long run. This is even on the SANS top 20 list.

To backup your data under Linux there are many solutions, all with various pro's and con's. There are also several industrial strength backup programs, the better ones support network backups which are a definite plus in a large non-homogenous environment. One of the other critical things to remember with backups is that whoever has access to them (backup

admin, cleaning staff) will have access to all your files unless you encrypt the backups. Physically securing backups is critical, damaging backups physically so they cannot be recovered is extremely easy, with magnetic media simply place a strong magnet near them, for CD's simply scratching the surface or cracking the CD will prevent usage. You should also keep a relatively recent set of backups offsite in case the building burns down or is inaccessible for some other reason (such as a chemical spill).

NON-COMMERCIAL BACKUP PROGRAMS FOR LINUX

There are numerous non commercial backup programs for Linux ranging from simple tools suitable for saving a few files to professional multi-system network backups.

TAR AND GZIP OR BZIP2

Oldies but still goldies, tar and gzip. Why? Because like vi you can darn near bet the farm on the fact that any UNIX system will have tar and gzip. They may be slow, klunky and starting to show their age, but it's a universal tool that will get the job done. I find with Linux the installation of a typical system takes 15-30 minutes depending on the speed of the network/cdrom, configuration another 5-15 (assuming I have backups or it is very simple) and data restoration takes as long as it takes (definitely not something you should rush). Good example: I recently backed up a server and then proceeded to blow the filesystem away (and remove 2 physical HD's that I no longer needed), I then installed Red Hat 5.2, and reconfigured all 3 network cards, Apache (for about 10 virtual sites), Bind and several other services in about 15 minutes. If I had done it from scratch it would have taken me several hours. Simply:

```
tar -cvf archive-name.tar dir1 dir2 dir3...
```

to create the tarball of all your favorite files (typically /etc, /var/spool/mail/, /var/log/, /home, and any other user/system data), followed by a:

```
gzip -9 archive-name.tar
```

to compress it as much as possible (granted harddrive space is cheaper than a politicians promise but compressing it makes it easier to move around). You might want to use bzip2, which is quite a bit better than gzip at compressing text, but it is quite a bit slower. I typically then make a copy of the archive on a remote server, either by ftping it or emailing it as an attachment if it's not too big (e.g. the backup of a typical firewall is around 100k or so of config files).

RSYNC

rsync is an ideal way to move data between servers. It is very efficient for maintaining large directory trees in synch (not real time mind you), and is relatively easy to configure and secure. rsync does not encrypt the data however so you should use something like SSH or IPsec if the data is sensitive (SSH is easiest, simply use "-e ssh"). rsync is [covered here](#).

AMANDA

Amanda is a client/server based network backup programs with support for most unices and Windows (via SAMBA). Amanda is BSD style licensed and available from: <http://www.amanda.org/>. Amanda now ships standard with a number of distributions. Commercial backup programs for Linux

BRU

BRU (Backup and Restore Utility), has been in the Linux world since as long as Linux Journal (they have had ads in there since the beginning as far as I can tell). This program affords a relatively complete set of tools in a nice unified format, with command line and a graphical front end (easy to automate in other words). It supports full, incremental and differential backups, as well as catalogs, and can write to a file or tape drive, basically a solid, simple, easy to use backup program. BRU is available at <http://www.tolisgroup.com/products3.html>.

QUICKSTART

Quickstart is more aimed at making an image of the system so that when the hard drive fails/etc. you can quickly re-image a blank disk and have a working system. It can also be used to 'master' a system and then load other systems quickly (as an alternative to say Red Hat's KickStart). It's reasonably priced as well and garnered a good revue in Linux Journal (Nov 1998, page 50). You can get it at: <http://www.tolisgroup.com/products3.html>.

Backup Professional

http://www.unitrends.com/br_bp.html

CTAR

<http://www.unitrends.com/ctar.html>

CTAR:NET

http://www.unitrends.com/br_ct.html

PC ParaChute

http://www.unitrends.com/ps_cr.html

Legato Networker

Legato Networker is another enterprise class backup program, now completely supported on Linux as both client and server. You can get it from: <http://www.legato.com/>.

BACKUP MEDIA

There are more things to back data up onto than you can drive a range rover over but here are some of the more popular/sane alternatives:

Name of Media	Pro's	Con's
Hard Drive	It's fast. It's cheap. It's huge (160 gigs). It's pretty reliable. (\$2-\$3 USD per gig)	It might not be big enough, and they do fail, usually at the worst possible time. Harder to take offsite as well. RAID is a viable option though.

CDROM	Not susceptible to EMP, and everyone in the developed world has a CDROM drive. Media is also pretty sturdy and cheap (\$0.20 USD per 650 Megs or so)	CDROM's do have a finite shelf life of 5-15 years, and not all recordables are equal. Keep away from sunlight, and make sure you have a CDROM drive that will read them.
Tape	It's reliable, you can buy BIG tapes, tape carousels and tape robots, but they're not very cheap.	Magnetic media, finite life span and some tapes can be easily damaged (you get what you pay for), also make sure the tapes can be read on other tape drives (in case the server burns down...).
Floppies	I'm not kidding, there are rumors some people still use these to backup data.	It's a floppy. They go bad and are very small. Great for config files though.
Zip Disks	I have yet to damage one, nor have my cats. They hold 100 megs which is good enough for most single user machines.	Not everyone has a zip drive, and they are magnetic media. The IDE and SCSI models are passably fast, but the parallel port models are abysmally slow. Watch out for the click of death.
Jazz Drives	1 or 2 gig removable hard drives, my SCSI one averages 5 meg/sec writes.	They die. I'm on my third drive. The platters also have a habit of going south if used heavily. And they aren't cheap. They are junk.
LS120	120 Megs, and cheap, gaining in popularity (hah, I actually believed that). These things are dead as far as I can tell.	Slow. I'm not kidding. 120 megs over a floppy controller to something that is advertised as "up to 3-4 times faster than a floppy drive".
Printer	Very long shelf life. requires a standard Mark 1 human being as a reading device. Handy for showing consultants and as reference material. Cannot be easily altered.	You want to retype a 4000 entry password file? OCR is another option as well.

AUTHENTICATION

Authentication is typically one of the two main lines of defense that systems and networks rely upon, so ensuring that your authentication subsystems are implemented correctly is important. The majority of Linux systems rely on usernames and passwords, while support for tokens, smartcards and other authentication systems are available they are still relatively rare. On top of this sits PAM, as far as I know all major vendors use PAM by default, so understanding how PAM works and using it correctly is very important.

PAM

"Pluggable Authentication Modules for Linux is a suite of shared libraries that enable the local system administrator to choose how applications authenticate users." Straight from the PAM documentation, I don't think I could have said it any better. But what does this actually mean? For example; take the program "login", when a user connects to a tty (via a serial port or over the network) a program answers the call (getty for serial lines, telnet or SSH for network connections) and starts up a login program, login then typically requests a username, followed by a password, which it checks against the /etc/passwd file. This is all fine and dandy until you have a spiffy new digital card authentication system and want to use it. Well you will have to recompile login (and any other apps that will do authentication via the new method) so they support the new system. As you can imagine this is quite laborious and prone to errors.

PAM introduces a layer of middleware between the application and the actual authentication mechanism. Once a program is PAM'ified, any authentication methods PAM supports will be usable by the program. In addition to this PAM can handle account, and session data which is something normal authentication mechanisms don't do very well. For example using PAM you can easily disallow login access by normal users between 6pm and 6am, and when they do login you can have them authenticate via a retinal scanner. By default Red Hat systems are PAM aware, and newer versions of Debian are as well (see below for a table of PAM'ified systems). Thus on a system with PAM support all I have to do to implement shadow passwords is convert the password and group files; and possibly add one or two lines to some PAM config files (if they weren't already added). Essentially, PAM gives you a great deal of flexibility when handling user authentication, and will support other features in the future such as digital signatures with the only requirement being a PAM module or two to handle it. This kind of flexibility will be required if Linux is to be an enterprise-class operating system. Distributions that do not ship as "PAM-aware" can be made so but it requires a lot of effort (you must recompile all your programs with PAM support, install PAM, etc), it is probably easier to switch straight to a PAM'ified distribution if this will be a requirement. PAM usually comes with complete documentation, and if you are looking for a good overview you go <http://www.sun.com/software/solaris/pam/>.

Other benefits of a PAM aware system is that you can now make use of an NT domain to do your user authentication, meaning you can tie Linux workstations into an existing Microsoft based network without having to say buy NIS / NIS+ for NT and go through the hassle of installing that. As far as I know all modern Linux distributions have PAM support and default to it.

PAM CRYPTOCARD MODULE

A PAM cryptocard module is available <http://projects.jdimedia.nl/index.phtml?ID=crypto&L=&BROW=1&W=1260&H=886>. Cryptocards are excellent for securing interactive logins since they do not require any special equipment on the client end, thus you can log in from a cybercafe for example with no fear of your password being stolen (since it changes each time you log in). Unfortunately Cryptocards tend to be expensive and require some user training, I would advise them primarily for installations with a higher need of security than "normal" or for infrastructure related servers and equipment (i.e. Authentication servers).

PAM SMART CARD MODULE

Smartcards can be used to sign and encrypt email as well as providing login services. The primary problem with smartcards however is that the client station needs a compatible card reader, the chances of finding these on a system outside of your office are slim indeed. A module to provide PAM support for smartcards is available <http://www.linuxnet.com/apps.html>.

PAM MODULE FOR SMB

SMB (Server Message Block) is incredibly popular protocol for the simple reason Microsoft has chosen to use it as their primary protocol for Windows 9x and NT (it is also supported in 2000). Many sites have existing NT infrastructures, adding Linux servers that require their own authentications infrastructure can be quite troublesome. Fortunately you can authenticate on Linux machines against SMB servers, packages are available http://rpmfind.net/linux/rpm2html/search.php?query=pam_smb and the primary site is http://www.csn.ul.ie/~airlied/pam_smb/. You can also install SAMBA on the machine and use this to authenticate but for workstations the PAM module is much more appropriate.

AUTHENTICATION SERVICES

Authentication services such as NIS and Kerberos are covered in the network servers section of the LAGS [here](#). Generally speaking they are easy to implement client side on modern Linux distributions, during install you are often given the choice of Kerberos, LDAP or NIS+ passwords and their related settings. Setting up the servers however is another matter.

PASSWORDS

In all UNIX-like operating systems there are several

constants, and one of them is the file /etc/passwd and how it works. For user authentication to work properly you need (minimally) some sort of file(s) with UID to username mappings, GID to groupname mappings, passwords for the users, and other misc. info. The problem with this is that everyone needs access to the passwd file, every time you do an ls it gets checked, so how do you store all those passwords safely, yet keep them world readable? For many years the solution has been quite simple and effective, simply hash the passwords, and store the hash, when a user needs to authenticate take the password they enter it, hash it, and if it matches then it was obviously the same password. The problem with this is that computing power has grown enormously and I can now take a copy of your passwd file, and try to brute force it open in a reasonable amount of time (assuming you use a poor hash system, or weak passwords).

USE A BETTER HASH

Using a hash such as MD5 or blowfish significantly increases the amount of computing power needed to execute a brute force attack, but there are two large problems with switching from the traditional crypt hash. The first is compatibility, if you use NIS or NIS+ with systems such as Solaris using a different hash then crypt will break authentication, obviously a problem. The other problem is that no matter how strong a hash you use poor passwords (such as the username or "dog") will still be easily discovered. If possible you should use a better hash, but if this is not possible then there is another solution.

USE SHADOW PASSWORDS

User account data is stored in /etc/passwd traditionally, but the actual password hashes and related data (password expiry, etc.) is stored in /etc/shadow, a file only readable by root. Programs that need to check a password can either run as root or use a setuid or setgid wrapper program (like PAM provides) to check the password, the only way to get access to /etc/shadow requires root privileges. There have been problems in past with setuid programs that read /etc/shadow leaking information, however these are relatively rare (and you are no worse off then storing passwords in a world readable location).

Several OS's take the first solution, Linux has implemented the second for quite a while now. Because most vendors rely on PAM for authentication services, implementing a new authentication scheme is relatively simple, all you need to do it add a PAM module that understands the new authentication scheme and edit the PAM config file for whichever program (say login) uses it. Now for an attacker to look at the hashed passwords they must go to quite a bit more effort then simply copying the /etc/passwd file.

CRACKING PASSWORDS

In Linux the passwords are stored in a hashed format, however this does not make them irretrievable, chances are you cannot reverse engineer the password from the resulting hash, however you can hash a list of words and compare them. If the results match then

you have found the password (the chances of a different word hashing to the same value as another are slim), this is why good passwords are critical, and dictionary based words are a terrible idea. Even with a shadow passwords file the passwords are still accessible by the root user, and if you have improperly written scripts or programs that run as root (say a www based CGI script) the password file may be retrieved by attackers. The majority of current password cracking software also allows running on multiple hosts in parallel to speed things up.

Most modern Linux distributions use MD5 hashed passwords at a minimum (notable exceptions are SuSE and Debian which default to crypt for backwards compatibility with NIS and the like). In any event password crackers will usually catch poor passwords or dictionary based passwords quickly. As well on modern systems passwords are protected in shadow password files, if an attacker has access to this file chances are they have sufficient privilege to do other things to compromise the system.

VCU

VCU (Velocity Cracking Utilities) is a windows based programs to aid in cracking passwords, "VCU attempts to make the cracking of passwords a simple task for computer users of any experience level.". You can download it http://packetstormsecurity.org/groups/wiltered_fire/NEW/vcu/

PASSWORD STORAGE

This is something many people don't think about much. How can you securely store passwords? The most obvious method is to memorize them, this however has it's drawbacks, if you administer 30 different sites you generally want to have 30 different passwords, and a good password is 8+ characters in length and generally not the easiest thing to remember. This leads to many people using the same passwords on several systems (come on, admit it). One of the easiest methods is to write passwords down. This is usually a BIG NO-NO; you'd be surprised what people find lying around, and what they find if they are looking for it. A better option is to store passwords in an encrypted format, usually electronically on your computer or palm pilot, this way you only have to remember one password to unlock the rest which you can then use. Something as simple as PGP or GnuPG can be used to accomplish this. If you can afford it using authentication tokens or smartcards are a good way to reduce the number of passwords you must memorize.

Many of these programs have been found to contain flaws, I advise using them with caution.

STRIP

Strip is a palm pilot program for storing passwords securely and can also be used to generate passwords. It is GNU licensed and available <http://www.zetetic.net/products.html>. The generation function is flawed and should not be used.

TO BE CONTINUED...

Editors Note:

The section, "filesystems" has been transposed from the original document due to space considerations. This transposition does not affect the meaning or effectiveness of the advice given in this guide.

2003 And Beyond

Author: Andrew Grygus aax@aaxnet.com

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

INTRODUCTION

Big corporate executives expect to be off pillaging some other company in just months, but you, the small business manager, need to plan 5, even 10 years out - because you'll likely still be stuck in the same business.

Most small business information systems weren't planned - they just got plugged in to do specific jobs - but in today's increasingly competitive markets, that isn't enough. Carefully planning an information strategy is critical when your competition is worldwide.

This article is a guide to trends that are already in full motion and well known by technology specialists, but are far from obvious to most business managers. I can't tell you what to do about them, without studying your particular business, but it will cast some light on what you should be looking at.

Much of this article deals specifically with Microsoft and Microsoft's future. This is inescapable, because Microsoft is a huge part of the information industry - and aspires to being all of it.

Throughout this work are numerous references to substantial articles from major on-line and print sources, so you can see I'm not just blowing smoke - I'm talking about real trends well known in the industry.

THE INFORMATION TECHNOLOGY INDUSTRY

The Technology Industry is currently in a deep slump, and will come back slowly and very selectively. PC manufacturers and most PC software publishers will not be among those recovering.

Corporate technology spending is expected to increase only marginally during 2003. A good part of that increase will be consumed by increased license costs for Microsoft products and will do little for the rest of the industry.

Much of the recovery will be led by the critical small business market (the bulk of business and employment in the U.S.), and that's a problem for big technology companies. Every time the corporate market slumps, the big guys declare their new small business initiatives, which always fall flat and are abandoned the moment the corporate market starts to recover.

Small businesses are adopting technology rapidly, by necessity, but products and services for small businesses become ever more diverse and more customized in response to intense competition. That means companies that successfully serve small business must themselves be small and very flexible. This does not favor the well-known industry leaders.

A big factor dragging down many technology companies is the increasingly tight relationship between computers and the telecommunications industry, as even voice starts being carried as Internet traffic (VoIP - Voice over Internet Protocol). The entire telecom industry, from top to bottom, is saturated with fraud, corruption, monopoly politics, heavy debt from unwise mergers and excess capacity. All this contributed to the current collapse, and will continue to affect big name network equipment makers like Lucent and Nortel, and even Cisco, through 2003.

INVESTMENT IN TECHNOLOGY

Technology stock prices will stay depressed through 2003, and will not approach the highs of the Dot.com boom any time soon - not until the next tulip craze comes along (oh, maybe 6 or 8 years). There will be solid, though highly selective, investment opportunities, along with significant risks, but for most investors, consumer industries will be less trouble and a lot safer.

Wireless data will be a fast growing field, but with few pure plays. Wireless equipment will mostly be side lines by the same network equipment companies wired network equipment is made by, and you know they're depressed for a while.

Keep in mind that "Telecom = Fraud", so avoid anything that has to do with telephones and phone lines unless you think you're sharp enough to play with the big guys who get paid to invent and perpetrate those frauds. Better to invest in wildcat oil wells

Internet Technologies, "Web Services", are another major growth area, but this field pretty much already belongs to Microsoft, IBM, Sun Microsystems, and other giants, as far as the corporate world is concerned.

Technology investors may start finding that participating directly in selected small businesses provides much better return than trading on the stock market, and will be a lot more productive for the economy. This is just about the only way to take advantage of the small business technology market.

EMPLOYMENT

The technology employment picture has been completely transformed by the Internet. The large number of corporate jobs that used to absorb entry level tech workers are being exported to India, Russia, Poland, and other places with high education and low pay - and those jobs aren't coming back. The bigger the company, the more jobs they'll be exporting. Most manufacturing and assembly work has already been sent to the Orient.

Highly skilled jobs remain, and pay well, because basic design, prototyping and pilot production will still be done here, but even if you have skills, there are problems getting hired. Human Resources departments haven't one clue in Hell how to evaluate skilled technical workers. They try to match exact experience and training to exact job requirements, and demand 3 years experience for specialties that have only existed for 18 months. Thousands of technical jobs go unfilled, while thousands who could do those jobs are asking, "Do you want fries with that?".

If you're over 40, the technology employment picture becomes really grim. You have abilities far beyond a 28 year old's, but he's the one that's going to get hired. For over 40s, there are many opportunities in consulting, especially in small and medium business, but that takes social skills and above all, selling skills, exactly what many chose technical careers to avoid.

So what does all this mean to the small business manager? Simple - whether you're looking for a consultant or an employee, you have a large pool of experienced over 40 technology workers to choose from, and you won't have to compete tooth and nail with corporate employers. Look for a track record of flexibility, and let the corps pay for the young ones' experience.

Oh, if you do want to pursue a technical career - one word: XML. Actually, a bunch of words, because to use XML effectively, you are going to have to learn a lot about business processes and methods too, and that means talking to a lot of non technical people.

THE PC INDUSTRY

The PC industry was once thriving, driven by rapid innovation. It's now down and it's not coming back. Microsoft's monopoly enforces the "Uniform Windows Experience" to the extent innovation and product differentiation to attract new computer purchases are simply impossible, even compared to nameless "white box" products.

IBM wisely withdrew from the general desktop PC market, while Hewlett Packard and Compaq merged out of desperation. Dell alone thrives by being just an assembler, well tuned to commodity markets. The other majors are weighted down with design engineers and manufacturing facilities. The few remaining second tier vendors like Gateway are struggling and

may fold, as Everex, AST, ALR, Packard Bell and others did in the last big shake-out.

All the major PC brands have moved to low end servers as their profit center, but the same "commoditization" is happening there. They all use the same Intel chips (because Windows won't run on anything else) and the same designs, and they all paint their boxes black. Essentially, they compete on price, which guarantees ever thinner profits. Meanwhile, Sun and IBM manufacture high end (non Windows) servers, which remain profitable.

"White Box" PCs by local builders are now the leading brand, having captured at least half the desktop market worldwide (I16), and now cutting into the low end server market. Declared dead by industry pundits a few years ago "White Box" has continued to grow its market share.

Why such success for the "no names"? To the "Local Brand" builder, PCs aren't a product, they're just part of a service offering, and that service is far beyond what the "Name Brands" could hope to offer. The quality is often higher, product consistency better, and if something does fail, it's usually fixed within hours with minimum hassle, often less time than you'd spend on the phone with HP or Dell.

With little possibility of regaining White Box territory, and unable to differentiate their products due to the "uniform Windows experience", name brand makers are locked in a pricing battle. Ever lower margins have forced them to cut heavily into customer support, and even Dell's once legendary support has been shaved to near worthlessness(I1). This makes them even more vulnerable to high support White Box products.

Competition for the business market may become even more intense, because I expect Microsoft branded "PCs" to take much of home market within a couple of years. There's already XBox, Home Gateway coming soon, and XBox 2 will be much more like a fully functional PC. A lot of other Microsoft branded hardware is appearing. Microsoft will use "security" and "Digital Rights Management" to force competitors out of the home market.

What does this all mean to the small business manager? It means the market has been turned upside down. If you want cheap, buy major brand computers. If you need support and effective service, buy your computers from a local service provider who provides White Box PCs and servers as part of his service offering.

Disclaimer: Automation Access builds servers and PCs as part of our services for customers, but market figures are from major industry publications.

PERIPHERALS

The peripherals business, led by printers, is in better shape than the PC industry, but there are significant downside risks, especially for HP (formerly Hewlett-Packard (I11)).

Dell, described by Sun's McNealy as a "grocery store that sells bananas as fast as it can", is master of commodity markets. Having already pushed PC giant HP into an ill-considered attempt to imitate Dell's marketing methods, now Dell is after the printer market too. They've signed a major deal with Lexmark to supply Dell branded printers and consumables. HP is very dependent on its profitable "imaging" business unit for income.

The threat to HP is not so much printers as consumables (ink and toner cartridges). Volume market printers are already sold at a loss, and the money is all made on overpriced consumables. Dell has stated they will be competing on the price of consumables (14). Lexmark, the #2 printer manufacturer won't be impacted as much, as it will "make it up in volume" as Dell's major supplier. Lexmark already manufactures HP compatible toner cartridges, so the competition may go beyond just the Dell printer / consumables value package.

What does this mean to the small business manager? It means careful selection of products may soon yield significant savings in TCO (Total Cost of Ownership) of printers when consumables are taken into account.

THE SOFTWARE INDUSTRY

The PC software industry is in the final days of being destroyed by Microsoft. Having leveraged a monopoly it was handed by IBM into multiple monopolies, with complete control over the PC manufacturers, and with an "Ethics? We've heard of it" attitude, Microsoft is preparing to drive the few remaining significant software publishers out of the Windows market.

Soon there will be Microsoft, Intuit, and Symantec. While Intuit will put up a strong fight, its popularity is not something Microsoft will tolerate for long. Revenue plans for **Microsoft Great Plains** do not allow for the existence of accounting software competitors. Microsoft will use Longhorn and .NET to bash and batter Intuit (see below). Symantec will continue because someone has to publish antivirus software, and it isn't going to be Microsoft (liability issues).

A few years ago, venture capitalists wouldn't fund a software startup if its product wasn't for Windows. Now they won't fund anything that runs on Windows either. Making a big splash and going public gave way to making a splash and selling out to Microsoft, but now Microsoft just tells companies, "Sell at our price or die", leaving little room for return on investment,

Ironically, the impossibility of commercial software competing against Microsoft's monopoly has spurred rapid growth in non-commercial software under the banner of **open source**. While much open source software runs on Windows, the majority runs on **Linux**, an operating system that is itself open source, thus available for free.

Linux has already taken big bites out of Microsoft's server sales and has blocked Windows from key accounts. Research firm IDC (the same firm Microsoft hired to "prove" Linux cost more than Windows) expects Microsoft's server market share to start to **decline** by the end of this year. Even more serious is the Linux threat to Microsoft's desktop monopoly, which becomes more credible by the month (15).

Paradoxically, a strong open source alternative is the best hope for a revived commercial software industry. Much software needed by businesses is simply of no interest to open source developers. As Linux becomes a mainstream business operating system, the market for commercial software running on Linux expands.

The market for commercial software running on Linux is, however, a market for small companies to serve, and will not spawn a "new Microsoft". **Microsoft has only two profitable products**, Windows and Microsoft Office. open source equivalents of both are already free (Linux and OpenOffice) or very inexpensive (commercial Linux and StarOffice).

Microsoft sees the danger very clearly, and is already spending millions of dollars to fight open source, but their most effective weapons are ineffective against such low cost products that are available from many sources, none of which are dependent on Microsoft for their success.

What does all this mean to the small business manager? It means you need to start preparing now for the "All Microsoft" future - or start preparing now to move away from the Windows environment. This decision may look to you like a "no brainer" right now, but it shouldn't look at all easy by time you've finished this article

"THE INTERNET"

The Internet is transport for a number of completely separate services (ftp, email, gopher, newsgroups, chat, dns, etc.), but to businesspeople, the Internet is **"the Web"** (World Wide Web service). Actually, they do use other services, but these are now integrated into the Web browser, so they appear to be part of the www service.

Web pages are rapidly becoming the essential **master communications center** for every business. For most businesses, the Web site will never produce income, directly, but will be critically important to generating revenue through established business methods.

A recent **Sears study** (17) confirmed for retailers what the Internet industry has known for some time. Customers who use Sears' appliance Web site generally travel to stores to purchase, but when they get there it is to buy, not to shop. They've already used the Web site to research products, compare value and make buying decisions. I'm sure you can see the potential for reducing sales staff this implies.

This trend is solid - people want to buy in the presence of the physical product, or talk to a real sales person on the telephone, **but people don't have time to go to stores to shop** any more. They can't afford to travel to stores except to buy. Shopping, research and price comparison can be done on the Internet quickly, conveniently, and at any time of day or night.

What does this mean to you, the small business manager? Without an effective, well designed and well promoted Web site, your business, products or services, **aren't going to be considered at all**. If people can't effectively gather the information they need before visiting your store or picking up the phone to acquire your services, they aren't going to bother, because they can get all that information **from your competitors**.

This bodes well for consultants, hosting services, and others who prepare and maintain commercial Web sites and Internet strategies, especially those that service small and medium business - in the long term. What's holding it back in the short term is that so many business people have yet to realize the critical importance of the Web site, and many more think they can do a fine job themselves in a couple of hours with Microsoft Front Page.

Increasingly, if the first impression made by a business' Web site is wanting, that will be the **last impression that business will get to make**. Speed and ease of finding information count for everything, and a slow Web site, or a Web site that offers pizzazz instead of content, will be the death of many businesses. Evolution is a harsh and unforgiving process.

At the other end, once hugely profitable Web consulting firms serving large business will continue their decline. Large businesses can now hire most of the expertise they need in-house at more attractive rates. There's plenty of "Dot.bomb" refugees desperate for jobs.

Profits go to companies that proceed carefully and with due consideration to all traditional business practices. Amazon has lost 2 billion selling books, while Powell's Books (I2) has been profitable all along.

Yes, there will still be room for "Web only" retail businesses, and perhaps Amazon.com will someday recoup its massive losses. Right now, Web sales are working best for highly specialized products and services with a very scattered clientele. If **S&M gear** is what you sell (I12), for instance, e-commerce could be your best bet for reaching your customers (it took me just a minute to find these WartenbergWheel guys).

The Internet will also be very strong in B2B (Business to Business) commerce, but not the way investors originally presumed. All the high profile customer / product / vendor match making sites and clearing houses went face to pavement, hard, and huge amounts of venture capital evaporated without a trace.

Finding the lowest price vendor turned out to be unimportant compared to serving **established business relationships**.

The Internet will increasingly replace current **EDI** (Electronic Data Interchange) non-standards, handling document interchange between established business partners using **XML** and other Internet syntax standards.

Once again, things look good for modest sized consultants, integrators and contractors who can do highly customized work, but there will also be major opportunities for big consultancies like EDS and IBM Global Services. In-house staff can't do your customers' or vendors' systems, so substantial third parties will be essential to integrate larger businesses

INTERNET TECHNOLOGIES

Not long ago, there were three network domains: **LAN** (Local Area Network), **WAN** (Wide Area Network) and "The Internet" (a very large peer-to-peer WAN). Today, there is only one: "**The Internet**". WAN traffic has become largely **VPN** (Virtual Private Network), encrypted traffic carried over Internet transport. LANs have become just local Internet subnets.

Protocols that were superior for LAN traffic, such as Novell's IPX/SPX have been dropped in favor of the Internet protocol, TCP/IP (except in organizations that want a very sharp demarcation between LAN and Internet for security reasons). The good news is that protocols that were inferior for any traffic, such as Microsoft's NetBEUI, are also being pushed aside.

Given this situation, we can talk of Internet technologies as if they applied equally to the LAN, because they do, even if the LAN isn't actually attached to the Internet (though that is becoming rare).

Not yet under monopoly control, Internet technologies are a seething hotbed of innovation - with one exception. Once Microsoft's Internet Explorer reached 70% market share, the once blazing hot evolution of Web browsers came to a sudden screeching halt. The only features being added to Internet Explorer now are features that allow Microsoft greater remote control of your network and computers.

The biggest things going on in Internet technologies are **Web Services** and **XML**. Both of these will deeply change the way your business does business in the future.

WEB SERVICES - XML, JAVA AND .NET

Web Services is becoming a very big thing, with the bulk of new system development within the corporate environment being rapidly moved to Web Services (I10).

Web Services is a system where programs run on Application Servers and are accessed by client Web browsers. Parts of the program may download and run on the local client, but the main program stays running on the server, takes requests from the clients and hands back results.

Web services clients may be PCs, but don't need to be, they just need to support a standards compliant Web browser. Communications is by XML, SOAP and Java. Whether you're running Windows, or OS/2, or OS/386, Linux, PalmOS, or what have you, is irrelevant (in theory).

You can access your Web services environment from anywhere you can get a network connection, and from any standards compliant platform. That network connection can be on the local network, or across the Internet (using a VPN, we hope, for security).

While doing any particular task, a client may be accessing Web services from several servers simultaneously, and application servers can even access services on other servers. This makes programming very modular - stuff needs only be written once and needs run only in one place to be accessed everywhere. This makes maintenance, upgrading and control very easy.

Web services can also be offered over the Internet by ASPs (Application Service Providers) for a fee, and these services will appear to be local to the client. Most Web services will be hosted on a company's own high speed network for best performance.

XML (eXtensible Mark-up Language) (I13), a critical key to Web Services, is an open standard syntax by which communication channels between very different computer systems can be established for exchange of documents containing data. For instance, a PC based accounting system could transmit an XML document that described a purchase order to a mainframe system, which could read and interpret the document and enter the information directly into its own order entry system.

XML is similar in principle to the HTML (HyperText Mark-up Language) so familiar for presenting and displaying Web pages, but XML is designed to handle data instead of page presentation. Web Services applications will take requests by XML document and return XML documents containing the requested response. Many expect XML to eventually replace EDI (Electronic Data Interchange).

XML is not a magic bullet that puts programmers out of work. It allows dissimilar computer systems to communicate with much greater ease than they traditionally have, but it does nothing to simplify the business logic within the system at either end. XML schemas have to be developed and interfaced with the business logic.

Numerous trade groups are developing XML schemas specific to the unique needs of the industries they

serve. Interfacing these schemas to the business logic of individual participating companies will be done on a case by case basis by programmers, it'll just be a lot easier than with EDI, and use the Internet as a transport (in most cases, but proprietary networks can still be used with XML).

Java (I14) is a programming language well suited to implementing Web Services, and is actively promoted by three of the major proponents of Web Services, **IBM**, **Bea Systems**, and **Sun Microsystems**. Java programmers are in high demand now, and the language is being used especially on the server side of the Web Services equation.

Interaction among very dissimilar systems is a major goal of Web Services and XML. Java is available for all significant computing systems, so strongly supports this goal.

Obviously, Java, XML and open standards provide too much freedom, flexibility and local control to be at all acceptable to the fourth major proponent of Web Services. We'll cover Microsoft's .NET Initiative (pronounced "DotNet") in detail in the **Microsoft** section. For now, .NET services may talk to other systems (yet to be demonstrated), but development must be done under Windows using mostly Microsoft tools, and .NET based services will run only on Windows.NET servers.

WIRELESS NETWORKING

Wireless networking is the hottest item in the technology press today, and is touted as the savior that will raise technology markets (especially hardware) from the dead. The implication is that if your business is not converting to wireless, it's falling way behind. The truth is a bit different.

Wireless has been around for some time, but only recently has it been able to overcome its traditional value proposition, "We're a lot more expensive, but we're a lot slower". Today, the price and performance differential is still there, but wireless is fast enough and cheap enough for a variety of uses. You use wireless where it is impractical to use wire, and nowhere where it is practical to use wire.

Even those magazines pushing wireless most heavily are replete with articles on wireless security, or more accurately, on the lack of **wireless security**. Securing a wireless network is possible, but more difficult, and it tends to go insecure if not monitored very carefully. Wireless has already been ordered pulled out of places it has been installed due to security problems, both in business and government.

In many businesses, wireless links are being installed by users for their own convenience, often to connect a laptop computer or to avoid the hassle of stringing a cable to the next room. In their competition to make installation easy, manufacturers ship their product with all security features turned off.

The proliferation of uncontrolled wireless access points has spawned the practice of **War Driving**, cruising the streets with simple detection equipment and marking buildings with graffiti indicating open access points for "free" Internet access. Industrial spies and thieves are doing the same, but without the graffiti.

Wireless in the home is growing quickly in popularity, because homes are relatively difficult to wire neatly, and the security risks are lower. Most home wireless systems are for shared Internet access, and with broadband access methods all less than 1.5-MB per second, low cost 11-MB/sec wireless networking is more than adequate.

MOBILE DEVICES

Accelerating the application of Wireless Networking is the rapidly rising sophistication and falling cost of mobile devices, ranging from cell phones up to Microsoft's "Tablet PC".

Microsoft will be pushing the Tablet PC very hard because its similarity to a full function PC means they can expect a monopoly position with these devices. Tablet PCs will probably find strong early adoption in medical facilities where users move around a lot, but are never far from the electrical power needed to recharge the devices.

Short battery life is the bane of devices that approach full Windows PC function, and also plagues smaller PDA (Personal Digital Assistant) devices running Microsoft's Pocket PC version of Windows. Battery life, and the shortcomings of handwriting recognition, will limit general acceptance of the Tablet PC.

Opposite the move from PC down to Tablet PC is the expansion of cell phones upward, adding text and photographic features as well as computing, data communication and Internet access. Microsoft has entered this market with its SmartPhone version of Windows (formerly code named "Stinger" (and popularly known as "Stinker")), but it faces very tough going.

Most handset manufacturers are highly suspicious of Microsoft's intentions, and those that aren't are idiots, as Sendo has proven so well. Motorola has selected Linux and Java for its phones (A10), while Nokia and Sony Ericsson have selected Symbian and Java (A11, A12), other devices will use Palm OS and Java. Due to misbehavior, Microsoft is forbidden by court order from implementing Java.

Designed from the start for small size and non-mouse operation, Opera is the leading Web browser for Internet access on mobile devices (A13). Microsoft, unable to compete using its bloated Internet Explorer, attempts to sabotage Opera where it can (R11). Opera's entertaining response to the latest attack has become a public relations classic (R12).

Because Microsoft will not be able to gain a dominant

position and crush innovation as it has done elsewhere, this will be a very dynamic field, allowing innovative companies both large and small to prosper, but there are dangers to be avoided.

A large number of VARs (Value Added Resellers) will develop specialized applications around devices using Microsoft mobile software, because that's an easy sell to clients so familiar with Windows. These VARs will stay small and specialized as a survival tactic - if their product's success results in a wide market, Microsoft will copy it and put them out of business.

Microsoft has become an extremely dangerous mobile integration partner, since it has acquired major accounting (Microsoft Great Plains), POS (Point of Sale - SMS QuickSell) and CRM (Customer Relations Management (B9) products, and no longer depends on third party business software products for integration. More ambitious (and smarter) VARs will select alternative environments and learn to deal with the selling issues.

Large scale mobile device integration by major vendors will increasingly deploy the Linux operating system, which is not only royalty-free and more flexible, but gives the developer complete control of their own future.

SECURITY

Since 9/11, network and computer security has become the hot topic of discussion. Other than huge amounts of discussion and thousands of column inches by "experts", formation of high level government committees and papers issued by them, nothing substantial has been done.

Further, nothing substantial is going to be done until there is a major disaster. When that happens, that disaster will be fixed and little more. Columnist Peter Stephenson relates the situation to Y2K (18). So much money was spent fixing Y2K problems that no disaster actually happened, so business executives feel they were ripped off (some of my clients weren't fixed in time and learned firsthand how serious the problem could have been). Now they think Network Security is just the latest rip-off.

In truth, the problems are very real, and will result in serious economic loss, and ruin for some businesses. Major avenues for disaster are:

Worms, Viruses and Trojans prey on the extreme vulnerability and seamless integration of the Windows monoculture. Business has been lucky that the fast moving worms launched to date have been unarmed, but that luck won't hold forever. There is no defense for the first few days after release of a new worm.

Home workers and business executives have access to company networks over secured access links, but their home computers aren't secure, leaving wide open holes through which crackers and industrial spies can

access business networks and through which passwords can be stolen. Microsoft's corporate network was raided by Russians this way.

Wireless Networking is a severe security problem for business networks everywhere. Employees are adding wireless links to corporate networks for their own convenience, and small businesses owners are using wireless to save the trouble of pulling cable. Manufacturers of wireless gear compete for ease of installation, so they turn security features off by default.

"War Drivers" are cruising around identifying businesses with insecure wireless links, and even spray paint symbols on the buildings indicating how to access the network. Industrial spies are doing the same, but without the graffiti.

Web Services are being deployed before security is worked out, providing many new avenues for invasion. In particular, Microsoft's .NET initiative aims for the same level of tight integration Windows has (it locks out competitors and promotes "ease of use"). This will result in similar security problems - once you're in, all the resources of the system are yours to exploit.

"Social Engineering" made Kevin Mitnick the most famous cracker of all, even though his technical skills were admittedly not that great. His new book, "The Art of Deception" (I9) shows both the extent and "ease of use" of this weakness, and presents suggestions on how to mitigate the problem.

File Sharing systems like Kazaa, Morpheus and the now defunct Napster, not only bog down your network, but expose your systems to worms, viruses and "spyware", create avenues for outsiders to steal sensitive files, and expose your business to multi-million dollar lawsuits from the entertainment industry.

Once the specialty of a very few highly skilled hackers, data theft and destruction is now an entertaining pastime for thousands of unskilled **"script kiddies"** and a standard tool for **industrial spies**, saboteurs, and even **law enforcement**. Losses will continue to climb out of control.

Spending on data system security will increase as companies try mitigating sharply increasing losses from worms, viruses, espionage, crackers and data theft, but this spending **will be ineffective**, because top management won't commit to the resources and sweeping changes required to secure their data systems effectively. They'll instead buy "band aid" fixes with big promises and little effectiveness.

The potential for massive destruction is shown by the recent spread of the W32/Klez.H worm (which is still raging out of control). Had Klez.H carried a destructive payload timed to go off about 15 days after release, the economic damages would have been stunning, and many business would have failed. Instead, it's **only cost \$9 billion** to clean up (so far), because the Klez.H

perpetrator disarmed it before release.

The potential for espionage is exemplified by the recent conviction of a prominent loan shark based on evidence from a "trojan" program (X9). Such programs are easily available to the public, and conveniently email the information they collect to the perpetrator (X10).

Don't think not being connected to the Internet will protect you either. Multi-gigabyte storage devices with the size and appearance of a key fob are now economically available, and plug into any convenient USB port. Not only can your critical data walk out on a ring of jingling keys, spy programs can be easily injected by anyone with brief access to your internal network.

While all systems have vulnerabilities, as recently demonstrated by the Slapper Worm (X42), the vast breadth and extent of vulnerability is a **uniquely Microsoft problem**, and for very specific reasons, which we'll cover in detail in the Microsoft section.

THE MICROSOFT "ROAD AHEAD"

With over 90% of the desktop PC operating system and office suite markets, "proven in court" monopolies yielding 85% and 79% profit margins (R2), \$43 billion in the bank, the highest capitalization (stock value) of any company, and having just been let off a major antitrust conviction with no punishment whatever, you might think Microsoft would be worry free.

Actually, Microsoft has something of a siege mentality, and that mindset is increasingly justified as the company fights too many battles on too many fronts. Expansion into new markets is meeting heavy resistance, and the old monopolies face unexpected threats. Open-source software and Internet protocols are leading concerns, as iterated by departing executive, David Stutz (R10).

In a recent shareholder meeting, Microsoft executives stated the company would pay **no dividend** from its \$43 billion because it still faced legal and other challenges, and needed all that money to defend itself (R3). In actuality, Microsoft faces many very real challenges of such magnitude that \$43 Billion may be a bit thin.

Other threats proved greater, however, so two months later, Microsoft **announced a dividend**. It's just 16 cents / share (\$850 million) (R6), a token dividend, but it's a dividend just the same, and another step on the path from "high growth" to "granny stock".

Why the dividend? There appear to be several reasons. For one, it served as a distraction from poor earnings projections, and the fact that Windows Desktop (Client) sales actually **fell by \$10 million** vs. the same quarter in 2001 (R8). Microsoft Office (Knowledge Worker) sales rose, but only by 8%, far from the 20%+ growth Microsoft is noted for.

Some have suggested the dividend was to get Bill Gates another \$100 million per year, which will be tax sheltered if Bush's "stimulus" plan prevails, but this seems an unlikely consideration for so major a policy change,

Most important was probably Microsoft's **stock price**, which has **trended down for three years**. Paying even a tiny dividend allows funds that require dividends to buy Microsoft shares for the first time. Microsoft hopes more buyers will bring the stock price up. Countering this is the growing feeling among investors that Microsoft is badly overvalued at its current price of 25 times earnings (R17). Only rapid growth can sustain this ratio, and that growth isn't happening.

At the same time, Microsoft announced a **2 for 1 stock split**. Now this looks really strange. Companies normally split their stock on the way up to make it more affordable by the share. Why is Microsoft splitting a stagnant stock? Apparently it's a desperation move to make the stock look cheaper, so "bottom feeders" will buy and move the price higher.

At the root of Microsoft's problems is its financial structure. **Leveraged by stock options** and other financial tricks (R7), it depends heavily on rapid revenue growth and increasing stock value. When you've saturated your market (over 90%), and that market is stagnant, rapid revenue growth becomes difficult. Should future growth look poor, holders of stock options are likely to cash out, and much of that \$43 Billion evaporates.

By necessity, Microsoft must expand into new markets. Problem is, they've never been particularly successful in entering any markets where they could not leverage their monopolies. As is true of other monopolies, they simply aren't competitive in open markets. While Windows has a profit margin of 85%, and Microsoft Office has a margin of 79%, **every other Microsoft division is losing money in reams** (R4).

Even where they can leverage their monopolies (corporate data centers, on-line services) Microsoft is meeting unexpectedly heavy resistance. Part results from serious weaknesses in Microsoft's products, part from battle hardened opponents now familiar with all Microsoft's tricks. In markets they already control, customers are becoming very **reluctant to upgrade** anything, because experience has proven it's a lot of expense for very little gain.

Another really big item is **open source software**. Open source products like **Linux** are rapidly locking Microsoft out of the midrange server market and are even becoming a threat to Microsoft's desktop monopoly, especially outside the U.S.. Microsoft's 10K filing with the SEC reflects this with warnings Microsoft could be forced to lower prices (R9).

Outside the PC arena, Microsoft faces entrenched adversaries with management far more astute and aggressive than anything they saw in the PC market. **XBox** for example, is taking heavy losses on each unit,

but still not selling enough of them to create the critical mass needed to generate interest among top game developers and gamers.

In the market for Interactive Television, Microsoft expected to dominate, and invested heavily in potential customers "just to make sure", yet they have been almost completely defeated there. The market has become dominated by companies like **Liberate**, **OpenTV**, and **TiVo**, with mostly Linux based products. Microsoft was unable to deliver an acceptable product, on time, and at competitive cost.

Microsoft's SmartPhone (Stinger) initiative to dominate the high end mobile phone market is on life support now that T-Mobile is said to be canceling or scaling back the program (R16) and Sendo is suing them for unfair business practices, misappropriation of intellectual property and just about everything else. Just about everyone else has signed up with Symbian and Java or Linux and Java.

To compensate, Microsoft is squeezing ever more revenue from current fully saturated markets by raising costs to their customers, mostly by changing licensing terms. This is creating resentment in formerly docile customers, many of whom consider the new terms **extortion**, and has them seriously looking at alternatives for the first time. It is also generating resistance to Microsoft's .NET initiative, now seen by many as a "vendor tie-in", ripe for exploitation.

Microsoft's public image, carefully crafted through billions (literally) spent with PR firms, continues to erode under the weight of license extortion, anti-trust action, license compliance raids, buying political influence, endless lawsuits over stolen products and patents, insecure and unstable products, obvious astroturf (fake grass roots) media campaigns, and much more

Innovation is another place where Microsoft is underperforming. They are fond of boasting about how much money they are pouring into R&D (Research and Development), yet not much new comes out of R&D. Some suspect the R&D efforts are financed simply to keep talented people away from startups and competitors, but it's probably simpler than that. Any innovation has to be retrofitted back onto the obsolete and hopelessly overcomplex Windows platform. This is highly limiting.

Some expect a combination of Microsoft's financial situation, market difficulties and erosion by open source software to result in an Enron style meltdown. This is highly unlikely. Microsoft management is too smart and not nearly corrupt enough to do an Enron, and they are coming up with many innovative ways to leverage their PC monopolies and squeeze more blood from any handy turnip.

Given a huge company that must grow rapidly, but which is encountering severe limits to growth, you can expect a lot of major changes in both products and policies, and that's what you are going to get.

LIMITING YOUR CHOICES

Making competing products unavailable is simply a quicker and more reliable method of helping you avoid "wrong" choices than developing a better product, and monopoly power gives Microsoft that option. Threatening software developers with no access to critical Windows information if they also programmed for OS/2 was typical - and concealed for years by an NDA (Non Disclosure Agreement). Expect the same technique to be used intensively against Linux.

Major PC manufacturers have always been under threat by Microsoft to eliminate products or configurations Redmond does not approve of - to "assure a **Uniform Windows Experience**". IBM's PC business, for instance, was severely damaged by Microsoft pressure to discontinue supporting OS/2. IBM got Windows 95 much later than other manufacturers, missing the introductory market, and paid more for it. **More recently**, Dell announced availability of Linux on many of its desktop PCs, but immediately withdraw the program without explanation.

Microsoft used this method to quickly reduce Netscape's browser share from 80%+ to nearly nothing. They were **convicted of multiple and very serious antitrust violations** in doing so, but the newly appointed Bush/Ashcroft Department of Justice declined to apply punishment or effective remedy. Microsoft is thus free to use similar methods to remove other products from the market.

Microsoft's current push is to have complete control over hardware design and availability. The practice of issuing joint Intel / Microsoft PC design specifications came to an end with the PC 2001 issue (J2). Microsoft alone now specifies PC design, leaving Intel as only a manufacturer (J3). This control is now made final and all encompassing by the Athens PC design.

Microsoft has three pressing reasons for seizing control of hardware design:

DRM (Digital Rights Management): Microsoft is assuring the motion picture and recording industries it will be the "safe" channel for distributing digital content that cannot be pirated. Credibility of this claim requires complete control of the hardware Windows will run on (see also Palladium and Home and Entertainment below).

Home Electronics: With XBox, Microsoft has started a move to control household electronics and entertainment. This goes hand in hand with their DRM efforts - Microsoft hardware will be the only means by which much entertainment content will be available, and competing DRM schemes can be blocked. The "Athens" PC extends this control to the telephone system.

Open Source: Competition from Open Source products

like Linux is being taken very seriously by Microsoft. It is already hindering their expansion in the business market, and they certainly don't want that to happen in the home market. By controlling the hardware platform they can assure that open source products do not have access to important hardware features. They have already stated that Palladium will not be ported to non-Windows platforms.

ATHENS

[Update 9 May 2003]: Microsoft has now specified the "next generation PC", code named "Athens" (R13) in precise detail. The specification ties PC design tightly to Microsoft's Windows initiatives and DRM (Digital Rights Management) plans, leaving no options for differentiation among PC vendors (R14). Athens is a "Microsoft Branded" PC in every detail except the label on the front.

Details of the specification raise questions as to whether any other operating system will be able run on this new machine, which will include Palladium chip based "security" features. As they now do with XBox (R15), Microsoft may be expected to pursue legal action against anyone who modifies Athens PCs to enable use of "unapproved" software.

Of particular concern are open source operating systems like Linux, which publish under the GPL (General Public License). The GPL forbids incorporation of code into a GPL'd product without releasing that code to the public, which would be forbidden by Microsoft's license terms. While the open source and server markets are already large enough to support their own hardware industry, "forking" the PC would raise costs for users and prevent use of these inexpensive machines in the consumer market, raising costs to consumers.

In addition, Athens is designed to **take control of your telephone communications**. Once again, Microsoft will use the highly successful tactic of distributing features to the lower tier users. Once users have deployed the features for their own convenience, they will force another uncontrolled Microsoft centric environment on management. [end update]

Microsoft does not limit its control to the hardware itself. Windows XP introduces a program of **"signed drivers"** and a tightly controlled **"Designed for Windows XP"** logo program (J4). If a hardware manufacturer does not submit drivers for approval by Microsoft (an expensive process), Windows XP pops up a warning message that the drivers are **not approved by Microsoft** and may cause problems with Windows.

Since Microsoft is the final authority that "signs" drivers and software packages (J5), it's clear they have the power to drag out the approval process if they don't favor a software or hardware product, or don't favor the company (perhaps because they offer Linux drivers too). Approval could even be denied entirely due to mysterious "incompatibilities".

The obvious weakness of these programs is that users can become accustomed to the unsigned driver warning, and products, both hardware and software, can still be sold saying something like "Works with Windows XP". It appears Microsoft is now moving to close this weakness and make unapproved products **entirely unavailable to the buying public**. Without sales, unapproved products will quickly disappear from the market, giving Microsoft complete control of what you can buy.

Office Depot has already issued a letter to suppliers informing them that products without the Microsoft "Designed for Windows XP" logo will no longer be carried by Office Depot as of May 2003 (J1, J6). Logically, this must be a response to pressure from Microsoft, and if other mass marketers do the same, that will effectively confirm it despite the nondisclosure agreements

Expect Microsoft to continue tightening these programs until unapproved products will simply not run at all on Windows. The justification for this, as with almost every anti-consumer and anti-competitive move Microsoft is making, will be to "enhance security".

These programs clearly tighten Microsoft's already powerful control over the availability of products that support competing environments, especially the availability of hardware drivers. Even years ago, with much less leverage, they were able to force Epson to drop printer driver support for OS/2.

The danger to Microsoft is that their heavy hand may spawn an alternative hardware industry, just as it spawned an alternative software industry. Large hardware companies are completely under Microsoft's control through dependency on Windows for volume, but new or smaller hardware companies, blocked economically from the Windows market, may choose to support alternatives.

This is most likely to happen in the relative safety of overseas locations, especially as overseas governments adopt Linux and open source. Here again, Microsoft's monopolies threaten to limit U.S. jobs, opportunities and American technology leadership.

LONGHORN - BIG CHANGES FOR WINDOWS

The very disappointing uptake of **Windows XP** has convinced Microsoft they must force upgrades. License 6 does force customers to upgrade on Microsoft's schedule, whether they want to or not, but a majority of the market has not adopted License 6, despite Microsoft's threats. Clearly drastic changes are needed to Windows to generate renewed upgrade revenue.

The successor to Windows XP (due in 2004, and rapidly slipping to 2005) is currently code named **Longhorn**, and it will not be compatible with your existing software, hardware or methods. Microsoft has already stated that backward compatibility will not be

a design feature.

Some expect the name Windows will be dropped completely. The antitrust agreement with the Bush DoJ specifically states "Microsoft Windows" throughout. By **maintaining incompatibility** (already planned due to design considerations), making it look different and calling it something else, Microsoft can free itself from antitrust oversight. "It's not Windows, it's a different product - the agreement doesn't apply."

The most important feature of Longhorn is replacement of the familiar DOS/Windows filesystem with an object database (W0). You will no longer copy files to a floppy or CD-ROM or attach them to an email, because there will be no files. Database records will be copied from one database to another, probably through a .NET server. Large organizations will have their own .NET servers, but everyone else will use one of Microsoft's, **a service for which you will pay a fee**.

The Longhorn filesystem will be based on the technology of a re-thought and expanded **SQL Server** database (the project coded Yukon) (W8). Obviously, SQL Server being so tightly integrated with the filesystem (W19) will have a negative impact on publishers of other database engines for Windows. Not strange then that market leaders **Oracle** and **IBM** are heavily pushing the Linux platform and barely mention their products run on Windows any more.

Current Windows based software will not be compatible with the Longhorn filesystem (W26). Microsoft has already stated that all their own software has to be rewritten for it - so will everyone else's. This will eliminate a huge number of software titles which are useful, but not sufficiently profitable to justify rewriting them. Others will fail because their conversion won't be done in time to compete with Microsoft products.

Coming with Longhorn is a new user interface, code named Sideshow (W5), so if you're currently trying to make sense of the new Windows XP user interface, 2005 is when Billy intends to yank your chain again. If you're a Windows programmer, you get to learn a new API framework named Avalon (W6) too.

Given Microsoft's enthusiasm for "rich data formats", I expect Longhorn is going to **eat disk space** at an alarming rate. Perhaps this is why Microsoft has suddenly taken a strong interest in storage technology and services (W9). It's also going to be a major backup problem, so watch for Microsoft to start offering .NET backup services (for which you will pay a fee). Others already offer ASP backup service, but expect incompatibilities with them that "assure security and data integrity" (unless, of course, they pay a large license fee to Microsoft).

I find it probable Longhorn will largely end the use of reliable, low cost servers (**Linux**, **NetWare**) for Windows users. This will set the stage for serious increases in licensing costs for already costly Windows server software.

Of course, Longhorn is going to be very late, so Microsoft is already hinting, then denying, they'll bridge the revenue gap with a Windows XP "Second Edition" upgrade (W10), or simply redefine "Longhorn" so they can get a partial product out in time. **License 6**, practically promises a major upgrade every 3 years or so, so I wouldn't even be surprised if they issued another all but unusable "end of the line" screw up like Windows Me.

The biggest risk to Microsoft is that the Longhorn effort falls apart, as did its "universal filesystem" predecessor, **Cairo** (W22), still an embarrassment to Microsoft. Cairo became later and later, was then "repositioned" as a "suite of technologies", and swept under the carpet. Failure of Longhorn would be more serious, because that would severely impact Microsoft's upgrade revenue stream.

Already, "Longhorn Server" has been dropped. The Longhorn filesystem will instead be an update to Windows Server 2003 (formerly Windows.NET Server 2003) and other features have been shoved down to "Blackcomb" (W21), the next scheduled Windows upgrade.

If you find all these Windows codenames confusing, a translation table is available

WINDOWS 2003

Windows 2003 has just been released [Update 9-May-03], and has brought with it some surprises, like the extent to which users will have to upgrade to run it. Win2003 is incompatible with all existing Microsoft server applications, except IIS 2003 (Internet Information Server) which was launched simultaneously with Windows 2003 (T9). IIS 2003 is not compatible with previous Windows versions, including Windows 2000 Server.

Patches are available to allow the Microsoft SQL Server database engine to run on Windows 2003 Server, but they are said to make the system very unstable (T9). SQL Server 2003 is still months away, and is a very expensive program, so upgrades are unlikely to be cheap.

Microsoft is already pushing Windows 2003 very hard to it's enterprise customers, pointing out that previous versions of Windows are dangerous to run because they're less stable, slower, and are riddled with serious security problems. That's the same thing they told us about Windows NT when 2000 came out, and the same thing they'll tell us about Windows 2003 when Windows 2006 comes out.

Complicating the situation is the fact that Microsoft doesn't consider Windows 2003 as released to be complete. Multiple features are to be released over a period of time (T11). These will have to be installed as available by users who wish to fully utilize this version of Windows

Arrayed against rapid upgrade are the compatibility problems, upgrade expenses, the poor state of the economy, and the fact that many enterprise customers are still in the middle of upgrading to Windows 2000. Analysts consider those still running on Windows NT 4.0 to be better candidates for immediate upgrading (T12). Of course, they're still on NT 4.0 because the upgrade to 2000/2003 is so traumatic.

A small but significant percentage of NT 4.0 / Windows 2000 customers have stated they will never upgrade to Windows 2003 because they are in the process of upgrading to Linux.

Microsoft has released benchmark showing Windows 2003 to be 69% to 89% faster than Linux for file and print services, but this "independent study" is highly suspect, comparing a highly tuned Windows system against an old version of Linux with known performance problems and an old version of Samba (T13). Both without optimization, using carefully orchestrated tests. At this point it is impossible to actually know if Windows 2003 is faster and under what circumstances it is.
[end update]

MICROSOFT OFFICE

Microsoft Office, not Windows is Microsoft's true cash cow, and it is Office, much more than Windows, that ties customers to Microsoft's expensive licensing plans. Currently, Office Pro costs about \$500 per computer, but some research firms expect it to reach \$700 per computer by the end of the decade. This will be combined with tough "one license per machine" enforcement measures.

Microsoft simply cannot allow inroads on Office's market share, yet the high and increasing cost of office, especially in view of License 6 (see below) has many businesses looking very hard at Sun Microsystems' low cost **StarOffice** and its "no cost" sibling, **OpenOffice**. Both run on **Windows** and **Linux**, and OpenOffice is being ported to Apple Macintosh. Both have excellent compatibility with Microsoft Office files, and native integration with XML.

Microsoft is also having a very hard time getting users to upgrade to the latest Office versions. Office 97 has long been replaced by Office 2000, which in turn has been replaced by Office XP, yet a huge number of Office users are still on Office 97 and show no signs of upgrading.

Microsoft's answer to both these problems is **Office 2003** (formerly named Office 11), currently in beta release and scheduled for final release in mid 2003. Office 2003 features a degree of tight integration with other Microsoft products that is impossible for other software vendors to achieve. It is also a degree of tightness Microsoft's customers will find nearly impossible to escape once committed. These features will be **required** by .NET and other Microsoft

initiatives.

Office 2003 and Windows Sever 2003 will include a **Rights Management Services** feature for document security (W25). If Microsoft can convince businesses to use this feature, Office 2003 documents will be completely unreadable by OpenOffice / StarOffice, WordPerfect Office, Lotus, and **by all older versions of Microsoft Office**, forcing a total upgrade of Windows, Office and the computers it runs on.

Update: A correspondent has told me that Rights Management interfaces for Office 2000 and Office XP can be downloaded from the Microsoft beta site. We will have to wait for formal release to see if these are available for the shipping product. This does not change the picture for Office 97 or for products that compete with Microsoft's.

Office 2003 will not run on Windows 95, 98, 98SE or Me. Microsoft is very clear that it will run only on Windows XP and Windows 2000 with SP3 (Service Pack 3) applied (W17). Currently over 60% of Microsoft's business customers are still running Windows 95/98, and would have to purchase all new computers for an XP upgrade - new computers soon to be obsoleted by Longhorn and Palladium.

Note that applying SP3 for Windows 2000 requires you to accept a license that allows Microsoft **to enter your computer systems, examine their contents and make changes without your knowledge or permission.** Some companies are refusing to apply SP3 even though it includes important security patches. The Windows XP license also includes these terms.

Next installment will cover new Microsoft technologies amongst other areas

(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>

AUUGN Book Review

Reviewed by Michael Still <mikal@stillhq.com>

GOOGLE™ HACKS: 100 INDUSTRIAL-STRENGTH TIPS AND TOOLS

BY TARA CALISHAIN AND RAELE DORNFEST. O'REILLY 2003
(ISBN 0-596-00447-8)

Did you know that there is a maximum of ten words in a Google™ search, or that repeating a word results in different search results? What about the fact that some search types are antisocial, and wont work with other search types? How about the fact that the order of your keywords affects the results that are returned?

Neither did I, and I use Google™ nearly every day.

I should probably start this review by letting you know that I don't believe that Google™ Hacks is not a book aimed at developers. It lists 100 interesting things you can do with Google™, in a quite accessible manner. That's not to say there is no code in Google™ Hacks - on the contrary, there are regular examples in Perl of how to parse saved Google™ output, and use the Google™ API - it is more that these examples are presented as utilities to use as is, and are not the topic of discussion in their own right. I would be comfortable with handing this book to a school child who needed to know how to use Google™ effectively. Especially if they had used the Internet before and know what a browser, and a URL are.

Then again, this doesn't mean that Google™ Hacks is simplistic. It still describes how to alter URLs, how to search using the Google™ special syntaxes, and so forth. Google™ Hacks even describes topics such as how to search date ranges (you can specify any Julian date range to Google™, although it must be an integer).

The Google™ API is a quite useful part of the Google™ offerings. The Google™ API is offered by via a SOAP interface, which makes it relatively easy to access for most programmers. Google™ Hacks describes how to reference the Google™ API in a variety of programming languages, and then goes on to show a series of example applications written in Perl.

Google™ Hacks also has a fun side. There is a section describing Google™ poetry, and Google™ art. The basic process of Google™ art is to post a message to a Usenet group (preferably one of the test groups) which contains keywords contrived in such a manner that a picture is displayed when a specific search is executed in Google™ Groups. The colour of each pixel results from the automatic keyword hi-lighting which occurs in Google™ Groups search results. Your picture is somewhat limited however, in that Google™ searches only support ten keywords, and therefore you can only have ten colours in your picture. The resolution supported by the text interface is also somewhat limited.

Some of the hacks described in Google™ Hacks are also quite non-hackish. For example, there is a whole section on how to build a useful website that Googlebot will find attractive to index. There is also discussion on possible formulations for the PageRank™ algorithm, as well as how to maximize Google's PageRank™ for your site.

Because each hack is meant to stand on it's own in this "reference mode" of use, there is a lot of repeated information. For example, many of the hacks describe how there is no access to a particular piece of functionality from the Google™ API, and that screen scraping in an automated manner is a breach of the Google™ terms of service. Similarly, the same prelude to setup the Google™ API is repeated in many of the examples of the API's use. Having read Google™ Hacks

from cover to cover, this repetition was really starting to get on my nerves by the end of the book.

In conclusion, once you bear in mind that Google™ Hacks is not aimed at developers, its a good read. It also makes an excellent reference for those who regularly use Google™, or are interested in learning more. The book is written in a manner which makes the vast majority of the content available to technical novices, whilst still discussing technical issues which really matter.

Linux and Open Source in Government 2004 CfP

Adelaide, 12–13 January 2004

Linux.conf.au is proud to announce the first annual Australian conference on Linux and Open Source in Government. Sponsored by Oracle and organised as a miniconference of the Linux.conf.au in cooperation with Linux Australia and AUUG, the conference will be held in Adelaide, South Australia, on 12 and 13 January 2004. Titled "The Challenges," this conference will focus on best practices, raise awareness and share experiences amongst Policy Makers and IT officials. Participants will come from Government and public departments, the academic sectors, and local, national and international organizations.

The Programme Committee invites proposals for papers for this conference. Suggested topics include, but are not limited to, the following:

- Demonstrations of Open Source Projects
- The Role of Open Source in Government
- State / Local Government using Open Source
- Citizen-Centric eGovernment
- IT Strategy and Enterprise Architecture
- Open Source Licenses and other Legal Issues
- Business Cases: Total Cost of Ownership
- Contemporary Security Issues
- Open Source and the Critical Information Infrastructure
- Open Source in the Military
- Open Source in Health Care and Bioinformatics
- Open Source Empowering People with Disabilities

Presentations may be given as 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.

SUBMISSION GUIDELINES

Those proposing to submit papers should submit an abstract (between ½ page and 2 pages), and a brief biography including:

- Full name (and preferred handle, if any)
- Email address

- Author's affiliation with commercial or relevant organisations
- Postal address
- Telephone and fax numbers, with area and country codes
- Short biography, in around 1–3 paragraphs.

If travel and accommodation assistance will be required, please list your closest international air terminal.

Abstracts and biographies should be submitted as plain text. Final presentations should be in the format of a 30–40 minute talk with 5–15 minutes for questions (a total of 45 minutes.) A written paper, for inclusion in the conference proceedings, must accompany all presentations. An electronic copy of your written paper, in an open format, is preferred over, well, a paper copy. We discourage the use of proprietary formats. Conference proceedings will be produced on A4-sized paper, and you should format your written paper accordingly.

Any papers that are accompanied by non-disclosure agreement forms will be rejected. All successful papers must be eligible for republication on-line and on distribution media given to conference attendees.

Linux.Conf.Au requires non-exclusive publication rights to accepted papers, including the publication of the audio proceedings as well as publication and reproduction rights to any video filmed during the presentations. Copyright ownership is retained by the author. You may be asked to sign an agreement to these conditions upon arrival at the conference.

Late submissions place an undue burden on our formatting team. In the event that you miss one of the deadlines we reserve the right to revoke any offer to present your paper.

At the submission stage only an abstract is required. Panel sessions will also be timetabled in the conference and speakers should indicate their willingness to participate, and may like to suggest panel topics.

SPEAKER INCENTIVE

Presenters of papers are afforded complimentary registration to this mini-conference.

IMPORTANT DATES

Deadline for abstracts/proposals	October 17, 2004
Notification of acceptance	October 31, 2004
Final submissions due	November 28, 2004
Conference	12–13 January 2004

For more information:
<http://lca2004.linux.org.au/ocgconf/>

AUUGN Book Review

Reviewed by Michael Still <mikal@stillhq.com>

ESSENTIAL CVS

BY JENNIFER VESPERMAN. O'REILLY 2003 (ISBN 0-596-00459-1)

I've been using CVS for quite some time now, but because of the way I work, and the sort of projects I work on, there are some features I have simply never had to use, or didn't even know existed. Essential CVS rounded out my knowledge of CVS well, and will form a valuable part of my reference library.

The structure of Essential CVS is logical. Essential CVS discusses what CVS is, why you would use it (including it's uses for systems administrators), and then immediately presents a quick start guide. The quick start guide takes the form of a cook book of the common operations you would perform in every day CVS use: importing projects; checking in and out files; committing changes; updating your sandbox; and so on.

Essential CVS then moves onto chapters which cover some of the specifics. For example, chapter two discusses every day use of CVS, in a little more detail than the quick start. As I read the book from cover to cover, this felt a little strange, as there was repeated content quite frequently. Then again, it is obvious that this part of the book is built as a reference, in which case repeating information makes the book much more user friendly.

The quick start chapter presents enough information for you to walk away and start using CVS. The later, more specific, chapters, then provide the information you need to solve more specific questions as you encounter them. The book takes detail and accuracy seriously, for example in several places it mentions that comments in the CVS source code imply that a given feature may change in the future, and user caution should be shown.

Essential CVS also provides a large volume of command reference information, listing the command line parameters to the various CVS tools, and their use. Based on a brief perusal of the CVS man page, this is probably a good thing, as I found the format of the command reference within Essential CVS to be much more readable than the man page.

As the majority of Essential CVS is a reference book, it live and dies by the quality of it's index. The difficult question as a reviewer is how to test an index? You could judge it on size, but that would only tell you that a lot of words were in the index. If those words are not relevant, then the index is still of little use. I chose instead to think of four or five things which I thought I might need with CVS, and then used the index to pull out the relevant information. Thus, if I could find the things I thought I needed in the index, then the index passes my simple test.

Essential CVS met all of my index tests. For all of the things I looked up in the index, which ranged from information on how to commit binary files sensibly to CVS, to how to access a repository remotely via SSH, and how to configure build scripts as a commit requirement, the index took me to the relevant information to within a page or two. Thus, I am happy to declare Essential CVS to be a useful reference book.

Jennifer is also an Australian, which is always a good thing. Essential CVS is a good reference book. If you're using CVS, planning on using CVS, or need a present for someone who does, then I heartily recommend this book.

AUUG Corporate Members

as at 1st June 2003

- ◆ ac3
- ◆ ANSTO
- ◆ ANU
- ◆ Apple Computer Australia Pty Ltd
- ◆ Aust Centre for Remote Sensing
- ◆ Australian Bureau of Statistics
- ◆ Australian Taxation Office
- ◆ Bradken
- ◆ British Aerospace Australia
- ◆ Bureau of Meteorology
- ◆ Cape Grim B.A.P.S.
- ◆ 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 Pty Ltd
- ◆ Deakin University
- ◆ Department of Land & Water Conservation
- ◆ Energex
- ◆ Everything Linux & Linux Help
- ◆ Fulcrum Consulting Group
- ◆ IBM
- ◆ Land and Property Information, NSW
- ◆ LPINSW
- ◆ Macquarie University
- ◆ Multibase WebAustralis Pty Ltd
- ◆ National Australia Bank
- ◆ NSW Public Works & Services, Information Services
- ◆ Peter Harding & Associates Pty Ltd
- ◆ Rinbina Pty Ltd
- ◆ St. Vincent's Private Hospital
- ◆ Sun Microsystems Australia
- ◆ 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

UNIX and Bioinformatics

Author: Vladimir Likic <likic@cybersource.com.au>

[Note: The author is a research scientist at the University of Melbourne, as well as a consultant in Bioinformatics]

We are now in the post-genomic era. Major advances in molecular biology and genome technologies have generated a series of scientific breakthroughs, many of which opened new frontiers of science. Genomes of dozens of different organisms have been sequenced, including that of *Mus musculus* (house mouse), *Drosophila melanogaster* (fruit fly), and *Homo sapiens* (human). A genome consists of stretches of DNA called genes, where genes are linear sequences of four bases (A, T, G, and C) whose combination conveys a specific message: a protein code. A protein molecule is a linear chain of amino acids (there are twenty commonly occurring amino acids). The central dogma of molecular biology is that DNA produces RNA which in turn is translated into proteins. Conceptually, the DNA sequence translates into a sequence of amino acids which makes a protein. In a living cell an elaborate "factory" is dedicated to making proteins from DNA blueprints. After a protein is produced, it spontaneously (or perhaps with a little help from other proteins) folds into an elaborate three-dimensional structure. The precise, spatial arrangement of amino acids as they appear in the three-dimensional structure of a given protein is absolutely essential for its biological function: extended, unfolded proteins are not functional.

Proteins are the fundamental building blocks of living organisms. The genome of a primitive bacterium *E. coli* encodes about 4000 distinct proteins. The human genome, which spells the blueprint of a human being, consists of about 3200 million DNA base pairs. It is estimated that human DNA contains about 30,000 genes, each of which encodes a protein with a specific biological function. To complicate matter further, alternative post-processing of the primary gene transcript can generate a variety of messenger RNAs coding for multiple "isoforms" of a given protein.

Knowing the genome's DNA sequence does not automatically translate into an understanding of the phenomenon of life. Having the genome sequence is akin to knowing the order of letters which make up a book written in an unknown language. In order to deduce the meaning conveyed by the book, one first needs to decipher the rules and the syntax of the underlying language. And by all means, the underlying language of life is extraordinarily complex. A protein sequence translates into its biological function in a complex manner, which also depends on the location of the protein within the cell, the amount of a given protein, the network of interactions that protein makes with other proteins, and a myriad of signals conveyed by simple messengers (such as calcium ions). Today the challenge is to understand these interactions as a whole. Techniques of modern high-throughput biology,

such as rapid sequencing, polymerase chain reaction, DNA microarrays, gene expression arrays, and others, are today widely employed to provide snapshots of these interactions in representative biological systems.

It is impossible to envision the great achievements of today's biology without massive data storage, processing power, and networking capabilities provided by modern computers. The central database of the National Center for Biotechnology Information (NCBI) in the USA known as GenBank started as a few hundred DNA sequences (back in those times NCBI employed technicians to type in the DNA code published in scientific journals on keyboards consisting of only four letters, A, T, G, and C). Today GenBank contains 22 million DNA sequences from over 130,000 different organisms, totaling 28,000 million base pairs of DNA! Due to advances in sequencing techniques in the past decade, the amount of newly generated data has been growing exponentially. In addition to databanks of DNA sequences, many other biological databanks have been established. These include databases of protein sequences and their annotations, protein three-dimensional structures, sequence motifs, protein expression patterns, metabolic pathways, and so on. Powerful computers are required not only for data storage and information retrieval, they are also needed for the analysis of data produced by techniques of modern biology. As evidenced by the growth of biological databases, the amount of biological data has soared in the past years.

What is the relationship between UNIX and bioinformatics? The field of bioinformatics relies on powerful computers for complex calculations, and servers which power networks and biological databases. With the decline of Digital and the once highly regarded VAX operating system in early nineties, today's general server market is split between UNIX and the OS from the company often referred to as the 900 pound gorilla with good salesman skills. One might expect that the same holds for the bioinformatics niche, which today commands billions of dollars in business. However, that is not the case. The bioinformatics niche is dominated by UNIX. And it is not a narrow advantage in market share that we are talking about: in the field of bioinformatics, UNIX rules.

A quick scan of major bioinformatics Web sites reveals that they are powered almost exclusively by UNIX variants. For example, at the time of this writing, NCBI (www.ncbi.nlm.nih.gov) runs Apache on Linux (which replaced IRIX sometime in 2002); DNA Data Bank of Japan (<http://www.ddbj.nig.ac.jp>) runs Apache on Solaris; The Wellcome Trust Sanger Institute (<http://www.sanger.ac.uk>), which played a key role in the sequencing of the human genome, runs Apache on Tru64; and so on. An educated guess is that most high-profile biological data banks run Oracle on UNIX, while smaller projects rely on open-source databases (MySQL or PostgreSQL) run on Linux or FreeBSD. But the influence of UNIX goes far beyond servers and databases. Most algorithms and programs which made

the bioinformatics revolution possible were originally developed on UNIX. And even today, many essential bioinformatics programs run only under UNIX. In short, a degree of familiarity with UNIX is essential for any serious bioinformatics endeavor. To give you an idea, nearly a third of the book "Developing Bioinformatics Computer Skills" by C.Gibas and P. Jambeck is dedicated solely to making the reader familiar with the UNIX environment.

In this article we review some well-known bioinformatics programs. One of the most fundamental tasks in bioinformatics is comparison of DNA or protein sequences. Sequence data are an extremely important and abundant type of biological data. Sequence comparisons are the basis for searching of sequence databases, identifying sequence motifs, inferring function of newly discovered proteins, building evolutionary trees, and many other tasks. We first examine FASTA and BLAST, two renowned programs for finding regions of local similarity between two protein or DNA sequences.

FASTA

The common problem in bioinformatics is to find all sequences from a database that are similar to a query sequence. The full mathematical solution to this problem is based on the methods of dynamic programming. While this solution is known, it is too complex for routine database searches. Although the computer speed increased many times in the past twenty years, the size of sequence databases increased even more! FASTA was one of the first programs widely used for searching of protein and DNA sequence databases. FASTA was developed by William R. Pearson and David J. Lipman (the paper describing FASTA appeared in the Proceedings of the National Academy of Sciences of the USA in 1988). FASTA employs a heuristic algorithm which allows fast searching of databases for similar sequences with excellent sensitivity. FASTA is maintained by William R. Pearson at the University of Virginia, and can be downloaded from <ftp://ftp.virginia.edu/pub/fasta/>. The FASTA package contains other useful programs, such as ALIGN which calculates a global alignment of two sequences, and LALIGN which finds the best local alignments between two sequences.

FASTA programs are distributed as the source code copyrighted by William R. Pearson and the University of Virginia. The distribution comes with several makefiles, and compiles easily under Linux. In the FASTA documentation W.R. Pearson comments "Over the years, as ATT Unix System 5 and BSD unix have converged these files have become very similar (..) I have tried to use very standard unix functions in these programs, and they have been successfully compiled, with very small changes to the Makefile, on Sun's (Sun OS 4.1), IBM RS/6000's (AIX), and MIPS machines (under the BSD environment)."

BLAST

BLAST (Basic Local Alignment Search Tool) refers to both a set of programs for the interrogation of

sequence databases, and the search algorithm used in those programs. BLAST is a heuristic algorithm for pairwise comparison of sequences developed to allow fast searching of large databases. BLAST programs solve the same problem as FASTA and they are nearly as sensitive, only faster. The paper describing BLAST was published in 1990 by S.F. Altschul, W. Gish, W. Miller, E.W. Myers, and D.J. Lipman.

The NCBI provides a free BLAST service for search against the latest version of GenBank (<http://www.ncbi.nlm.nih.gov/BLAST/>), and precompiled BLAST binaries are available for download (<ftp://ftp.ncbi.nih.gov/blast/>). Stand-alone BLAST binaries are provided for IRIX, Solaris (Sparc/Intel), DEC OSF1, Linux/Intel, HP-UX, MacOS X, and Win32. Stand-alone BLAST allows for BLAST search to be performed on a local machine, against locally stored databases. NCBI also provides binaries for the BLAST client (blastcl3) which can be used to access the NCBI BLAST search engine from the command line. For example, if the query sequence is stored in the file p230.fasta, the BLAST search against the GenBank protein ("blastp") non-redundant database ("nr") can be executed as:

```
$ blastcl3 -p blastp -d nr -i p230.fasta -o p230.blast
```

This of course assumes a connection to the Internet, and the results of BLAST search will be stored in the file p230.blast.

In addition to stand-alone BLAST and BLAST client, NCBI also provides the binaries for a standalone BLAST WWW server. The BLAST WWW server allows one to set up in-house BLAST Web service which can provide a search service against a local database. UNIX was clearly a major platform for the development of both the BLAST algorithm and BLAST programs. The people at NCBI supply win32 versions of BLAST, but those who need the BLAST WWW server for Windows are out of luck. As it is explained on the NCBI Web site: "At this time the Standalone WWW BLAST Server is only available for UNIX web servers".

HMMER

While FASTA and BLAST programs are intended for interrogation of a sequence databases against a single sequence, a much more sensitive database search can be done against a family of sequences. Proteins which have similar functions in different organisms often share regions of sequence similarities, and a simple logic dictates that the similar portions must be those that are important for the common function. Given a set of such proteins one could build a mathematical model which in some way describes these conserved regions, and the degree of sequence conservation in these regions. Such a model then could be used to interrogate a database of proteins with unknown functions in order to find out which of these may be homologous to the family. The interrogation is performed in a biased way, by heavily weighting similarities in conserved regions (because they are likely to be conserved in homologous proteins), and

giving little weight to the sequence outside conserved regions (because those can be variable without affecting the function). Depending on the quality of the model, such a biased search may be much more sensitive than plain FASTA or BLAST.

A very powerful method for building models of similarity regions within a family of proteins is based on the profile hidden Markov models (HMMs). A renowned implementation of profile HMM is the program HMMER, developed by Sean Eddy at the Washington University in St Louis. This entire problem with protein families can be turned around: if enough protein families are known (a family is sets of similar proteins with a common function) one could build a database of profile hidden HMMs and then a single sequence could be searched against such a database to reveal to which family the protein may belong! The most useful database of profile hidden HMMs is PFAM, which employs HMMER for the calculation of its HMMs. The July 2003 version of PFAM contains multiple sequence alignments and hidden Markov models for 6190 distinct protein families.

HMMER is available for download from the Washington University in St Louis (<http://hmmerr.wustl.edu/>). Unlike FASTA and BLAST, HMMER is released under the GPL. As first sight this may not seem as a big deal, especially to those used to Linux and GNU programs. However, the GPL license (or other truly free licenses) are not a rule in the field of bio-computing. Most programs distributed as the source code seem to be released under dual licenses: typically, these licenses allow free and unrestricted use of the program in an academic environment, but the use in a commercial environment is often associated with hefty fees. Thus the release of a renowned program such as HMMER under the GPL is a big deal. In fact, Sean Eddy and co-workers have placed the entire PFAM database under the GPL, which means that HMMER and PFAM will remain freely accessible to anyone who needs it.

HMMER has been developed and employed under UNIX. The source code for HMMER can be downloaded from the HMMER home page, and should compile cleanly on any UNIX platform (and indeed compiles so on Linux). HMMER developers also provide an impressive list of precompiled binaries for various UNIX variants: Tru65/Alpha, Linux/Alpha, Solaris/SPARC, Solaris/Intel, Linux/Intel, Linux/HP-IA64, HP-UX/HP-IA64, AIX, OpenBSD, FreeBSD, IRIX, and Apple OS/X! What about HMMER on Windows? To quote the installation documentation "[HMMER] should also compile on Microsoft Windows platforms, but you would have to work around the GNU configure script and UNIX makefiles". Ouch!

VMD/NAMD2

Molecular visualization programs can display protein three-dimensional structures on high resolution computer screens. One such program is VMD from the Theoretical and Computational Biophysics Group at University of Illinois at Urbana-Champaign led by Klaus Schulten. VMD allows one to display, animate, and analyze proteins and DNA structures using 3-D

graphics. VMD has also built-in Tcl and Python scripting languages, a feature which distinguishes serious from less advanced molecular visualization programs. VMD is distributed free of charge (with registration), together with the source code. As usual with programs which deal extensively with graphics, precompiled binaries can save a lot of time and frustration. The VMD home page (<http://www.ks.uiuc.edu/Research/vmd/>) offers a plethora of UNIX VMD binaries, including AIX, Solaris, HP-UX, IRIX, Tru64, Linux, and MacOS X. Binaries for Windows 2000/XP using OpenGL are also available. Molecular visualization is most often done on desktop computers. Having in mind that Windows holds lion's share of the desktop market, it is interesting that about 50% of VMD downloads are for UNIX variants. Not surprisingly, the latter are dominated by Linux and Mac OS X versions. For other UNIX binaries the number of downloads is much smaller. However, the VMD developer John Stone notes that "the number of (Unix) machines using VMD is much larger than the numbers indicated by the download stats, as it seems that sysadmin types typically download one copy of VMD then install it on a whole lab of Unix workstations."

Protein molecules exhibit continuous internal motion at physiological temperatures, known as "protein dynamics". These motions consist of a great variety of jiggling and wiggling, the details of which depend on protein's sequence and the three-dimensional structure. It is now well understood that protein internal motions are important for biological function. For example, many enzymes become inactive when cooled to low temperatures (which is a way to quench internal motions). A careful examination of the equilibrium three-dimensional structure of hemoglobin (the oxygen carrier protein in blood) shows that the oxygen molecule cannot reach the heme group which is buried deeply within the protein. The fact that hemoglobin does bind oxygen is because the structure of hemoglobin fluctuates, and these fluctuations leave just enough space for oxygen to enter deep into the protein core and bind to the heme group.

Protein dynamics is very difficult to study experimentally and a great deal of what is known today comes from computer simulations, known as molecular dynamics simulations. NAMD2 is a modern program for molecular dynamics simulations, and also comes from Klaus Schulten's group. In contrast to VMD which is about graphics, NAMD2 is about number crunching. As stated on its home page (<http://www.ks.uiuc.edu/Research/namd/>), "NAMD is a parallel, object-oriented molecular dynamics code designed for high-performance simulation of large biomolecular systems. NAMD scales to hundreds of processors on high-end parallel platforms and tens of processors on commodity clusters using switched fast ethernet". Similarly as VMD, NAMD2 is distributed free of charge with the source code, and precompiled binaries for more than a dozen UNIX variants are available for download. NAMD has been developed on UNIX, but the program has also been ported to Windows since version 2.2. The user survey for year

2000 showed the following distribution of operating systems among NAMD users: Linux (65%), IRIX (23%), Solaris (5%), AIX (4%), HP-UX (2%), Tru64 (1%), T3E (1 %).

MOLSCRIPT AND RASTER3D

As everyone in academia knows, getting the results is only half the job. The other half, often more difficult, is getting it published. Nowadays scientific publications require sophisticated illustrations. It is often said that a picture is worth a thousand words, but today's prestigious scientific journals expect pictures that are worth more than that. When someone solves a three-dimensional structure of a new protein, for example, it is expected that the research article describing the structure will contain the figure depicting "in a nutshell" the new structure. This may a high quality raster image showing individual atoms as balls, or something schematical, such as a cartoon-like drawing which shows only protein sub-structures such as alpha-helices and beta-strands.

MOLSCRIPT is an excellent program for the generation of publication-quality protein pictures, developed and maintained by Per Kraulis. To get an idea about the "art of molecular images" visit the MOLSCRIPT picture gallery

(<http://www.avatar.se/molscript/doc/molscript.html>). MOLSCRIPT is free for academic use, and is distributed as a source code. MOLSCRIPT has been developed under UNIX, and according to the documentation "The installation procedure has been tested (so far) only on an SGI IRIX system, but should work on other UNIX systems as well." MOLSCRIPT compiles easily on Linux. A quick search of the Internet reveals a few web pages which offer advice on how to compile MOLSCRIPT on Windows. One such page starts with "Unfortunately, yes - molscript requires bison to interpret command input and a compiler (with OpenGL support) in order to generate the binaries". Indeed, unfortunate news which spell trouble for those who want MOLSCRIPT on Windows.

RASTER3D is a set of programs for the preparation of high-quality images of molecular three-dimensional structures. While MOLSCRIPT is the best for artistic and schematic pictures, RASTER3D shines for photo realistic raster images (see the picture gallery at <http://www.bmsc.washington.edu/raster3d/raster3d.html>). RASTER3D was developed under UNIX, and it has been tested under Solaris, HP-UX, IRIX, Digital UNIX, and of course Linux. It is distributed as the source code. As for the Windows version, precompiled binaries are available from the RASTER3D home page with the following note: "Precompiled win32 binaries for 2.6c contributed by Suhaib Siddiqi. Installation instructions are here. Totally unsupported."

CONCLUSIONS

The programs described here were chosen from a rather subjective viewpoint, but with an aim to convey just how great the variety of tasks in modern bioinformatics can be. Other programs were considered for inclusion in this article (and would

certainly complete this picture), but were left out because of space limitations. Those include: EMBOSS, the European Molecular Biology Open Software Suite (<http://www.hgmp.mrc.ac.uk/Software/EMBOSS/>); MEME, for the discovery of motifs in protein sequences (<http://meme.sdsc.edu/meme/website/intro.html>); PHYLIP, for studies of evolutionary relatedness of organisms (<http://evolution.genetics.washington.edu/phylip.html>); MMTK, a Python based molecular modeling toolkit (<http://starship.python.net/crew/hinsen/MMTK/>); Sparky, NMR assignment program (<http://www.cgl.ucsf.edu/home/sparky/>); the BioPerl project (<http://www.bioperl.org/>); and so on. However, even if all these programs were described in some detail, many important tasks in modern bioinformatics would not be mentioned. For example, multiple sequence alignment, protein secondary structure prediction, homology modeling, gene finding, prediction of protein cellular location, and others.

For those interested in finding out more about bioinformatics, the book by C.Gibas and P. Jambeck "Developing Bioinformatics Computer Skills" (O'Reilly, 2001, ISBN: 1-56592-664-1) is a good place to start. For an introductory text on bioinformatics, "Introduction to Bioinformatics" by Arthur M. Lesk (Oxford University Press, 2002, ISBN 0-19-925196-7) is recommended. The book "Bioinformatics: A Practical Guide to the Analysis of Genes and Proteins" by Andreas D. Baxevas and B.F. Francis Ouellette is a classic in the field (Wiley-Interscience, 1998, ISBN 0-471-19196-5). Finally, the book "Statistical Methods in Bioinformatics: An Introduction" by Warren J. Ewens and Gregory R. Grant (Springer, 2001, ISBN 0-387-95229-2) covers statistics behind modern bioinformatics, is reserved for those who want to get down to nitty-gritty details of sequence comparison methods.

ACKNOWLEDGMENTS

Author thanks Trevor Lithgow, Emanuel Strehler, and Jamie Fletcher for reading this article while in preparation, and for valuable comments and corrections.

Tuning and Optimizing Red Hat Linux Advanced Server for Oracle9i Database

Author: Werner Puschitz <<http://www.werner.us>>

The following procedure is a step-by-step guide with tips and information for tuning and optimizing Red Hat Linux Advanced Server for Oracle9i.

This summary (HOWTO) shows how I tuned and optimized Red Hat AS 2.1 (kernel 2.4.9-e.3, 2.4.9-e.10smp; glibc 2.2.4-26, 2.2.4-29.1) for Oracle 9iR2

(9.2.0).

A procedure for installing Oracle9iR2 on Red Hat AS 2.1 can be found at

<http://www.puschitz.com/InstallingOracle9i.shtml>
<http://www.puschitz.com/OracleOnLinux.shtml>

INTRODUCTION

Please point out every error you can find. I welcome email from any readers with comments, suggestions, or corrections. My address is webmaster@puschitz.com. I will continue to update and add new information for this article. So make sure to come back. :)

Before you begin making any changes to the Linux systems, make sure that the Oracle database is down!

ORACLE LIMITS ON LINUX

Some limits apply to Red Hat Advanced Server only.

Linux supports 64-bit file I/O on 32-bit Intel platforms.

According to the white paper <<http://otn.oracle.com/tech/linux/pdf/9iR2-on-Linux-Tech-WP-Final.PDF>> "Oracle9iR2 on Linux: Performance, Reliability and Manageability Enhancements on Red Hat Linux Advanced Server 2.1", the limits are as follows:

- Number of files per database: 64K
- Number of blocks per file: 4 million
- Maximum block size: 16 KB
- Maximum size for a database file is 64 GB
- Maximum database size is 4 petabytes with 16 KB blocks

On a 4 GB RAM machine, the size of the SGA (SGA utilizes shared memory) can be increased up to is 2.7 GB. This requires changes in Linux and Oracle. By default, the maximum size is 1.7 GB.

On a 8 GB RAM machine, the size of the SGA can be increased up to 7 GB by using the shared memory filesystem <<http://lwn.net/2001/1206/a/tmpfs.php3>> "shmfs".

A maximum size of 5.4 GB of SGA can be created using the "bigpages" feature for System V shared memory where the page size is 4 MB vs. the regular 4 KB.

On a machine that supports Physical Address Extension (PAE), the SGA can theoretically have a size of 62 GB. The PAE mechanism allows addressing using 36 bits on IA-32 systems. But current hardware limitations and practical consideration limit the actual size of the SGA on such systems.

The number of local concurrent users on a 4 GB server in non-MTS mode can range from 600 through 1200 without becoming unacceptable slow. For more information on the tpcrc run that measured the number

of concurrent users, see <http://otn.oracle.com/tech/linux/pdf/1_linuxVM_v2_accepted.pdf> "Linux Virtual Memory in Red Hat Advanced Server 2.1 and Oracle's Memory Usage Characteristics"

WHY USE RED HAT ADVANCED SERVER

Red Hat Linux Advanced Server has several features and enhancements that don't exist in other Red Hat versions. Among other things, Red Hat AS provides:

- Asynchronous I/O
- Process scheduler with CPU affinity, cache affinity, and per CPU runqueues and locks that provide better performance
- "mapped base" (base address for shared libraries) can be changed dynamically allowing larger sizes for the SGA
- Page frame of size 4 MB as opposed to 4 KB can be used for the SGA which improves performance for large SGAs
- The kernel can also use the "high memory" pool (physical memory above 1 GB) for allocating page table entries (PTE) which allow a higher number of Oracle connections
- Elimination of copy to bounce buffer improves I/O performance

UPGRADING THE LINUX KERNEL

The recommended kernel for Red Hat Enterprise Linux 2.1 is 2.4.9-e.25. This kernel has several fixes that are relevant to Oracle including fixes for memory problems and *kswapd* problems.

If the Linux server has <= 4 GB RAM, the kernel "kernel-smp" should be used for SMP machines, or the kernel "kernel" should be used for UP machines. If the Linux server has > 4 GB RAM, the enterprise kernel "kernel-enterprise" should be used for UP and SMP machines.

To check if these kernels are installed, execute e.g. the following command:

```
rpm -q kernel-smp kernel-enterprise
```

To check which kernel is currently running, execute the following command:

```
uname -a
```

To install e.g. the enterprise kernel, download the "kernel-enterprise" RPM and execute the following command:

```
rpm -ivh kernel-enterprise-2.4.9-e.25.1686.rpm
```

To make sure that the right kernel is booted, check the */etc/grub.conf* file if you use GRUB, and change the "default" attribute if necessary. Here is an example:

INTRODUCTION (HEADING 3)

Payload text here...

```
default=1
timeout=10
splashimage=(hd0,1)boot/grub/splash.xpm.gz
title Red Hat Linux (2.4.9-e.25enterprise)
    root (hd0,1)
    kernel /boot/vmlinuz-2.4.9-e.25enterprise ro
root=/dev/hda2 hdc=ide-scsi
    initrd /boot/initrd-2.4.9-e.25enterprise.img
title Red Hat Linux Advanced Server (2.4.9-e.25smp)
    root (hd0,1)
    kernel /boot/vmlinuz-2.4.9-e.25smp ro
root=/dev/hda2 hdc=ide-scsi
    initrd /boot/initrd-2.4.9-e.25smp.img
title Red Hat Linux Advanced Server-up (2.4.9-e.25)
    root (hd0,1)
    kernel /boot/vmlinuz-2.4.9-e.25 ro
root=/dev/hda2 hdc=ide-scsi
    initrd /boot/initrd-2.4.9-e.25.img
```

In this example, the "default" attribute is set to "1" which means that the 2.4.9-e.25smp kernel will be booted. If the "default" attribute would be set to "0", then the 2.4.9-e.25enterprise kernel would be booted.

After you installed the new kernel and/or made changes to the `/etc/grub.conf` file, reboot the server.

Once you are sure you don't need the old kernel anymore, you can remove the old kernel by running:

```
su - root
rpm -e <OldKernelVersion>
```

When you remove the kernel, you don't need to make any changes to the `/etc/grub.conf` file.

NOTE: Be very careful when removing a kernel! Making a mistake could render the server unbootable.

SIZING SWAP SPACE

In order to perform a typical Oracle 9i installation and to create a simple prototype database, Oracle says that you need a minimum of 512MB of RAM for the Oracle9i Server, and the amount of swap space should be equal to twice the amount of RAM or at least 400 MB, whichever is greater. Oracle also says that the minimum swap space should be at least the same as physical memory size.

Swap Size Recommendations

To summarize Oracle's recommendation for the database and to take system configurations into account that were used for workload testings, here is what I came up with:

```
e 1 GB RAM      2 GB - 3 GB Swap Space
  2 GB RAM      2 GB - 3 GB Swap Space
  3 GB RAM      3 GB Swap Space
  4 GB RAM      4 GB Swap Space
  8 GB RAM      4 GB Swap Space
 16 GB RAM      8 GB Swap Space
```

The swap space will not be utilized until the system

runs out of physical memory. So don't configure too much swap space. Keep in mind that if the system starts using swap space, it has a negative impact to the performance of the database. So make sure that the system has always enough physical RAM and that it doesn't use swap space continuously.

Checking Physical Memory

You can check the size of physical memory by running the following command:

```
grep MemTotal /proc/meminfo
```

You can find a detailed description of the entries in `/proc/meminfo` at <http://www.redhat.com/advice/tips/meminfo.html>

Alternatively, you can use "free(1)" to check the memory:

```
# free
              total    used    free shared buffers  cached
Mem: 1031004  734656 296348      0 262404 287388
-/+ buffers/cache: 184864 846140
Swap: 2097144    40184 2056960
#
```

In this example the total amount of available memory is 1031004 bytes. 184864 bytes are used by programs and 846140 bytes are available for more programs.

Don't get confused with the first line that shows that 296348 bytes are free! If you look at the usage figures you can see that most of the increase of memory is for buffers and cache. Linux tries to use all the memory for disk buffers and cache. It helps the system to run faster because disk information is already in memory and Linux doesn't have to read it from disk again. If space is needed by a program or application like Oracle, Linux will make the space available immediately. So if your system runs for a while, you will usually see a small number for "free" in the first line, and there is nothing to be worried about.

Note: If you create a large SGA (shared memory) and start the database, "free" won't show all the memory that has been allocated for SGA as "used" right away. That's because Linux does not assign page frames to a memory mapping right after it has been created due to reasons of efficiency.

Checking Swap Space Size and Usage

You can check the size and current usage of swap space by running the following command:

```
cat /proc/swaps
```

If your swap partition is not large enough, you can add another swap partitions to your system. See <http://www.redhat.com/docs/manuals/linux/RHL-8.0-Manual/custom-guide/s1-swap-adding.html>

"Adding Swap Space" for more information. Adding a permanent swap file to the system is not recommended due to the performance impact of the filesystem layer.

SETTING SHARED MEMORY

Shared memory allows processes to access common structures and data by placing them in shared memory segments. It's the fastest form of IPC (Interprocess Communication) available since no kernel involvement occurs if data is passed between the processes.

Oracle uses shared memory segments for the SGA (Shared Global Area) which is an area of memory that is shared by all Oracle background and foreground processes. The size of the SGA has a major impact to Oracle's performance since it holds database buffer cache and much more.

To see all shared memory settings, run:

```
ipcs -lm
```

Setting SHMMAX Parameter

This parameter defines the maximum size in bytes for a shared memory segment. Since the SGA is comprised of shared memory, SHMMAX can potentially limit the size of the SGA. Ideally, SHMMAX should be large enough so that SGA can fit into one segment.

The default size on RH 2.1 AS is 33554432. With this value, the Oracle Database Configuration Assistant failed on my server with the following error message:

```
ORA-27123: unable to attach to shared memory segment
```

Setting SHMMAX to 1 GB always worked for me when I setup a medium sized database. However, it is suggested that it should be set to 2 GB; the default maximum size of the SGA is 1.7 GB which requires a larger SHMMAX. And if the available size of the SGA is set to 2.7 GB by changing "mapped base" at the Linux OS level, then SHMMAX should be set to 3 GB.

The maximum value of SHMMAX can be set to 4GB-1.

(A typical 32-bit Linux system without Physical Address Extension (PAE) is divided into 3 GB user space and 1 GB kernel space.)

The default shared memory limit for SHMMAX can be changed in the proc file system without reboot:

```
su - root
echo "2147483648" > /proc/sys/kernel/shmmax
```

Alternatively, you can use *sysctl(8)* to change it

```
sysctl -w kernel.shmmax=2147483648
```

To make the change permanent, add the following line to the file */etc/sysctl.conf*. This file is used during the boot process.

```
echo "kernel.shmmax=2147483648" >> /etc/sysctl.conf
```

Setting SHMMNI Parameter

This parameter sets the maximum number of shared

memory segments system wide. The default number on RH 2.1 AS is 4096. To my knowledge this value should be sufficient.

```
# cat /proc/sys/kernel/shmni
4096
```

Setting SHMALL Parameter

This parameter sets the total amount of shared memory in bytes that can be used at one time on the system. The default size on RH 2.1 AS is 2097152. To my knowledge this value should be sufficient.

```
# cat /proc/sys/kernel/shmall
2097152
```

SETTING SEMAPHORES

Semaphores can best be described as counters which are used to provide synchronization between processes or between threads within a process for shared resources like shared memories. System V semaphores support semaphore sets where each one is a counting semaphore. So when an application requests semaphores, the kernel releases them in "sets". The number of semaphores per set can be defined through the kernel parameter SEMMSL.

To see all semaphore settings, run:

```
ipcs -ls
```

The SEMMSL Parameter

This parameter defines the maximum number of semaphores per semaphore set.

Oracle recommends to set SEMMSL to the largest PROCESSES init.ora parameter of any database on the Linux system plus 10.

Oracle also recommends to set SEMMSL to a minimum value of 100.

The init.ora parameter PROCESSES specifies the maximum number of operating system processes that can be started by the Oracle instance. In a non MTS environment, Oracle spawns a system user process for each connection. This means that in such an environment the PROCESSES parameter defines the maximum number of simultaneous Oracle connections minus sum of all Oracle background processes. It can also be said that the PROCESSES value should never be greater than SEMMSL.

The SEMMNI Parameter

This parameter defines the maximum number of semaphore sets (identifiers) in the entire Linux system.

Oracle recommends to set SEMMNI to a minimum value of 100.

The SEMMNS Parameter

This parameter defines the total number of

semaphores (not semaphore set) in the entire Linux system. A semaphore set can have more than one semaphore, and according to the *semget(2)* man page, values greater than SEMMSL * SEMMNI makes it irrelevant.

Setting it to a minimum value of 256 is for initial Oracle installation only.

Oracle recommends to set SEMMNS to the sum of the PROCESSES parameter for each database on the system, adding the largest PROCESSES twice, and then adding 10 for each DB.

The maximum number of semaphores that can be allocated on a Linux system will be the lesser of:
SEMMNS or (SEMMSL * SEMMNI)

Setting SEMMSL and SEMMNI to 100 makes sure that SEMMNS semaphores can be allocated as determined by the above calculation.

The SEMOPM Parameter

This parameter defines the maximum number of semaphore operations that can be performed per *semop(2)* system call.

The *semop(2)* function provides the ability to do operations for multiple semaphores with one *semop(2)* system call. Since a semaphore set can have the maximum number of SEMMSL semaphores per semaphore set, it is often recommended to set SEMOPM equal to SEMMSL.

Oracle recommends to set SEMOPM to a minimum value of 100.

Setting the Semaphore Kernel Parameters

To determine the values of the four described semaphore parameters, run:

```
cat /proc/sys/kernel/sem
250 32000 32 128
```

Alternatively, you can run:

```
ipcs -ls
```

All four described semaphore parameters can be changed in the proc file system without reboot:

```
su - root
# echo SEMMSL value SEMMNS value SEMOPM value
SEMMNI value >/proc/sys/kernel/sem
# These are the values I'm using since I don't
# want to lower Red Hat's default values.
# The only value I raise is SEMOPM to comply with
# Oracle's minimum requirement for SEMOPM.
echo "250 32000 100 128" > /proc/sys/kernel/sem
```

Alternatively, you can use *sysctl(8)* to change it:

```
sysctl -w kernel.sem="250 32000 100 128"
```

To make the change permanent, add or change the following line in the file */etc/sysctl.conf*. This file is

used during the boot process.

```
echo "kernel.sem=\"250 32000 100 128\"" >>
/etc/sysctl.conf
```

To see the new updated semaphore settings, run:

```
ipcs -ls
```

SETTING FILE HANDLES

The maximum number of file handles denotes the maximum number of open files that you can have on the Linux system.

Setting System Wide Limit for File Handles

The value in */proc/sys/fs/file-max* sets the maximum number of file handles or open files that the Linux kernel will allocate. When you get error messages about running out of file handles, then you might want to raise this limit. The default value on RH 2.1AS is 8192.

For an Oracle server it is recommended that the file handles for the entire system is set to at least 65536.

To determine the number of file handles for the entire system, run:

```
cat /proc/sys/fs/file-max
```

The maximum number of file handles can be changed in the proc file system without reboot:

```
su - root echo "65536" > /proc/sys/fs/file-max
```

Alternatively, you can use *sysctl(8)* to change it:

```
sysctl -w fs.file-max=65536
```

To make the change permanent, add or change the following line in the file */etc/sysctl.conf*. This file is used during the boot process.

```
echo "fs.file-max=65536" >> /etc/sysctl.conf
```

Setting File Handles Limit for the Oracle User

There is still a per user limit of open files which is set to 1024 by default:

```
$ su - oracle
$ ulimit -n
1024
$
```

To change this, you have to edit the file */etc/security/limits.conf* as root and make the following changes or add the following lines, respectively:

```
oracle soft nofile 4096
oracle hard nofile 8192
```

The "soft limit" in the first line defines the number of

file handles or open files that the Oracle user will have after login. If the Oracle user gets error messages about running out of file handles, then the Oracle user can increase the number of file handles in this example up to 8192 ("hard limit") by running the following command:

```
ulimit -n 8192
```

You can set the "soft" and "hard" limits higher if necessary.

You also need to make sure that `pam_limits` is configured in the file `/etc/pam.d/system-auth`. This is the PAM module that will read the `/etc/security/limits.conf` file. The entry should read like:

```
session required /lib/security/pam_limits.so
```

Here are the two "session" entries I have in my `/etc/pam.d/system-auth` file:

```
session required /lib/security/pam_limits.so
session required /lib/security/pam_unix.so
```

Now login to the oracle account again since the changes will become effective for new login sessions only.

```
$ su - oracle</font>
$ ulimit -n
4096
$
```

The default limit for oracle is now 4096 and the oracle user can increase the number of file handles up to 8192:

```
$ su - oracle
$ ulimit -n
4096
$ ulimit -n 8192
$ ulimit -n
8192
$
```

SETTING ASYNCHRONOUS I/O

Red Hat Advanced Server supports asynchronous I/O in the kernel. Asynchronous I/O permits Oracle to continue processing after issuing I/Os requests which leads to much higher I/O throughputs. This enhancement also allows Oracle to issue thousands of simultaneous I/O requests with a single system call. It also reduces context switch overhead.

According to a Red Hat webcast I attended, only 2 Oracle dbwriter processes are needed when asynchronous I/O is being used.

To enable Oracle to use asynchronous I/O, it is necessary to relink Oracle. Oracle ships Oracle9iR2 with asynchronous I/O support disabled. According to Oracle, this is necessary to accommodate other Linux distributions that do not support asynchronous I/O.

Relinking Oracle to Enable Asynchronous I/O for

Oracle9iR2

```
# shutdown Oracle
SQL>; shutdown
```

```
$ su - oracle
$ cd $ORACLE_HOME/rdbms/lib
$ make -f ins_rdbms.mk async_on
$ make -f ins_rdbms.mk ioracle
```

The last step creates a new "oracle" executable "\$ORACLE_HOME/bin/oracle". It backs up the old oracle executable to \$ORACLE_HOME/bin/oracleO, it sets the correct privileges for the new Oracle executable "oracle", and moves the new executable "oracle" into the \$ORACLE_HOME/bin directory.

If asynchronous I/O needs to be disabled for any reason, run the following commands:

```
# shutdown Oracle
SQL> shutdown
$ su - oracle
$ cd $ORACLE_HOME/rdbms/lib
$ make -f ins_rdbms.mk async_off
$ make -f ins_rdbms.mk ioracle
```

Enabling Asynchronous I/O in init.ora for Raw Devices

The `disk_async_io` init.ora parameter needs to be set to true:

```
disk_async_io=true
```

Note that this init.ora parameter is already set to true by default:

```
SQL> select value, isdefault from v$parameter where name = 'disk_async_io';
```

VALUE	ISDEFAULT
TRUE	TRUE

Enabling Asynchronous I/O in init.ora for Filesystem Files

Make sure that all Oracle datafiles reside on filesystems that support asynchronous I/O (e.g. "ext2"). According to Oracle's white paper <<http://otn.oracle.com/tech/linux/pdf/9iR2-on-Linux-Tech-WP-Final.PDF>> "Oracle9iR2 on Linux: Performance, Reliability and Manageability Enhancements on Red Hat Linux Advanced Server 2.1", Oracle9iR2 has been certified with the standard Linux filesystem "ext2" on RH AS 2.1. In addition, Oracle has also been certified for raw devices.

The `disk_async_io` init.ora parameter needs to be set to true (same as for raw devices):

```
disk_async_io=true
```

Note that this init.ora parameter is already set to true by default:

```
SQL> select value, isdefault from v$parameter where
```

```
name = 'filesystemio_options';
VALUE          ISDEFAULT
-----
none          TRUE
```

The `filesystemio_options` init.ora parameter needs to be set to `asynch`:

```
filesystemio_options=asynch
```

This init.ora parameter is platform-specific. By default, this parameter is set to `none` for Linux and thus needs to be changed.

```
SQL> select value, isdefault from v$parameter where
name = 'filesystemio_options';
```

```
VALUE          ISDEFAULT
-----
none          TRUE
```

The `filesystemio_options` can have the following values with Oracle9iR2:

- `asynch`: This value enables asynchronous I/O on file system files.
- `Directio`: This value enables direct I/O on file system files.
- `Setall`: This value enables both asynchronous and direct I/O on file system files.
- `None`: This value disables both asynchronous and direct I/O on file system files.

Increasing I/O Throughput at the Linux OS Level

The `/proc/sys/fs/aio-max-size` parameter can be changed if asynchronous I/O is used for Oracle datafiles residing on filesystems (e.g. "ext2"). To my knowledge, this parameter does not have any effect to raw devices. According to the <http://otn.oracle.com/tech/linux/pdf/9iR2-on-Linux-Tech-WP-Final.PDF> "Oracle9iR2 on Linux: Performance, Reliability and Manageability Enhancements on Red Hat Linux Advanced Server 2.1" document, Oracle9iR2 has been certified with the standard Linux filesystem "ext2" on RH AS 2.1.

To get better I/O throughput for Decision Support Systems (DSS) workloads, the `/proc/sys/fs/aio-max-size` parameter should be increased to > 1 MB. A typical DSS system queries large amount of data and makes heavy use of full table scans. Parallel Query is particularly designed for DSS.

For Online Transaction Processing (OLTP) workloads, the default size of 131072 would suffice. A typical OLTP system has high throughputs, are insert- and update-intensive, have concurrent access by many users, and have large, continuously growing data volume.

To determine the number of bytes, run:

```
su - root
# cat /proc/sys/fs/aio-max-size
131072
```

The maximum number of bytes can be changed for e.g.

DSS systems in the `proc` file system without reboot:

```
echo "2147483648" > /proc/sys/fs/aio-max-size
```

Alternatively, you can use `sysctl(8)` to change it:

```
sysctl -w fs.aio-max-size=2147483648
```

To make the change permanent, add or change the following line in the file `/etc/sysctl.conf`. This file is used during the boot process.

```
echo "fs.aio-max-size=2147483648" >>
/etc/sysctl.conf
```

INCREASING SPACE FOR LARGER SGA (2.7 GB) TO FIT INTO MEMORY

If the size of SGA does not need to be increased from 1.7 GB to 2.7 GB, then the following steps can be skipped.

By default, the maximum size for SGA is 1.7 GB on a 32-bit system without Physical Address Extension (PAE). You will also be able to allocate 1.7 GB SGA if you have less than 4 GB RAM. In this case you have to make sure you have enough swap space, however, this will have an impact to the performance of the database. I was even able to bring up a database with a SGA size of 2.64 GB on a test PC that had 256 MB RAM.

Theoretically, the SGA can have a size of up to 62 GB on a system that supports Physical Address Extension (PAE). The PAE mechanism allows addressing using 36 bits on IA-32 systems. But current hardware limitations and practical consideration limit the actual size of the SGA on such a system. Since I do not have such a system, I will not cover the steps for creating SGAs larger than 2.7 GB via the "tmpfs" filesystem.

To increase the size of the SGA to 2.7 GB without using a shared memory filesystem (tmpfs), the following needs to be done:

- The base address "mapped base" for Oracle's shared libraries has to be lowered at the Linux OS level.
- Oracle needs to be relinked with a lower base address for SGA which uses shared memory segments.

Address Mappings on Linux - Shared Memory and Shared Library Mapping on Linux

Normally, the 4 GB address space for a 32-bit Linux system is split into 4 equal sized sections for different purposes:

```
0GB-1GB User space - Used for executable and
brk/sbrk allocations (malloc uses brk for small
chunks)
```

```
1GB-2GB User space - Used for mmap's (shared memory),
shared libraries and malloc uses mmap (malloc uses
mmap for large chunks).
```

```
2GB-3GB User space - Used for stack.
```

3GB-4GB Kernel Space - Used for the kernel itself.

- The *mmaps* grow bottom up and the stack grows top down. The unused space used by the one can be used by the other.
- The split between userspace and kernelspace can be changed by setting the kernel parameter *PAGE_OFFSET* and recompiling the kernel. By default, the *PAGE_OFFSET* macro yields the value *0xc0000000*.
- The split between *brk(2)* and *mmap(2)* can be changed by setting the kernel parameter *TASK_UNMAPPED_BASE* and recompiling the kernel. However, on Red Hat AS this parameter can be changed for individual processes dynamically without reboot or kernel recompilation.

Usually, the portion of address space available for mapping shared libraries and shared memory segments consists of virtual addresses in the range of *0x40000000* (1 GB) - *0xc0000000* (3 GB). On Red Hat AS, *0x40000000* is the default base address for shared libraries and shared memory segments. The default base address for mapping shared memory segments can be changed and overwritten for programs and applications by non-root users. The default base address "mapped base" for loading shared libraries for programs and applications can be changed by the user root only.

The default base address that Oracle uses for SGA (shared memory segment) is *0x50000000* and not *0x40000000*. Oracle uses or keeps the space from *0x40000000-0x50000000* for loading Oracle shared libraries. As I mentioned before, *0x40000000* is the default base address on RH AS for loading shared libraries which can only be changed by the user root. Oracle increased the base address for SGA to prevent address range conflicts between the segments (shared memory segment and shared libraries).

If the base address for shared memory segments would be *0x15000000* and if the base address for shared libraries would be *0x40000000*, then Oracle cannot create the SGA larger than *0x2b000000* bytes or 688 MB, even though there is address space available above the shared libraries portion. (According to Oracle, Oracle binaries will no longer work if the base address for shared memory segments is lower than the base address shared libraries like in this example. Even though I didn't experience any problems, I would not recommend it).

If the base address for shared memory segments is *0x50000000* and if the base address for shared libraries is *0x40000000*, then Oracle can create a SGA that starts at *0x50000000* and ends almost at *0xc0000000*; *0xc0000000* is the address where the kernel address space begins. This means that the SGA can have a size of almost *0x70000000* bytes or 1.792 GB - actually it's about 100 MB less due to stack space and other use of memory.

Once again, Oracle increased the default base address for SGA to *0x50000000* so that all shared libraries can be loaded below *0x50000000*, and the rest of the space up to almost *0xc0000000* can be used for shared

memory.

You can verify the address mappings of Oracle processes by viewing the proc file */proc/<pid>/maps* where *<pid>* stands for the Oracle process ID. The default mapping of an Oracle process might look like this:

```
08048000-0ab11000 r-xp 00000000 08:09 273078
/ora/product/9.2.0/bin/oracle
0ab11000-0ab99000 rw-p 02ac8000 08:09 273078
/ora/product/9.2.0/bin/oracle
0ab99000-0ad39000 rwxp 00000000 00:00 0
40000000-40016000 r-xp 00000000 08:01 16 /lib/ld-
2.2.4.so
40016000-40017000 rw-p 00015000 08:01 16 /lib/ld-
2.2.4.so
40017000-40018000 rw-p 00000000 00:00 0
40018000-40019000 r-xp 00000000 08:09 17935
/ora/product/9.2.0/lib/libodmd9.so
40019000-4001a000 rw-p 00000000 08:09 17935
/ora/product/9.2.0/lib/libodmd9.so
4001a000-4001c000 r-xp 00000000 08:09 16066
/ora/product/9.2.0/lib/libskgxp9.so
42606000-42607000 rw-p 00009000 08:01 50
/lib/libnss_files-2.2.4.so
50000000-50400000 rw-s 00000000 00:04 163842
/SYSV00000000 (deleted)
51000000-53000000 rw-s 00000000 00:04 196611
/SYSV00000000 (deleted)
53000000-55000000 rw-s 00000000 00:04 229380
/SYSV00000000 (deleted)
bffffb000-c0000000 rwxp fffffc000 00:00 0
```

As this address mapping shows, shared libraries start at base address *0x40000000*. The address mapping also shows that Oracle uses the base address *0x50000000* for SGA (in this example System V shared memory for SGA). Here is a summary of all the entries:

The text (code) section is mapped at *0x08048000*:

```
08048000-0ab11000 r-xp 00000000 08:09 273078
/ora/product/9.2.0/bin/oracle
```

The data section is mapped at *0x0ab11000*:

```
0ab11000-0ab99000 rw-p 02ac8000 08:09 273078
/ora/product/9.2.0/bin/oracle
```

The uninitialized data segment *.bss* is allocated at *0x0ab99000*:

```
0ab99000-0ad39000 rwxp 00000000 00:00 0
```

The base address for shared libraries is *0x40000000*:

```
40000000-40016000 r-xp 00000000 08:01 16 /lib/ld-
2.2.4.so
```

The base address for SGA (System V shared memory) is *0x50000000*:

```
50000000-50400000 rw-s 00000000 00:04 163842
/SYSV00000000 (deleted)
```

The stack is allocated at *0xbffffb000*:

```
bffffb000-c0000000 rwxp fffffc000 00:00 0
```

Now it should become clear what needs to be done to

provide more space for SGA. To increase the space for SGA, two base addresses need to be changed. The base address "mapped base" for shared libraries needs to be lowered at the Linux OS level, and the base address for SGA (shared memory) needs to be lowered at the Oracle level (application level).

Note: Once the base addresses have been changed at the Linux OS level and at the Oracle level, all Oracle commands need to be executed with a lower "mapped base"! This means that every new shell must run with a lowered "mapped base". Further down I will show you how you can automate this so that every Oracle user gets automatically a shell with a lowered "mapped base".

Changing the Base Address "mapped base" for Shared Libraries at the Linux OS Level

The default base address "mapped base" on RH 2.1AS is TASK_UNMAPPED_BASE = 0x40000000 (decimal 1073741824 or 1 GB). This is the address that splits the section between *brk(2)* and *mmap(2)*, which defines available space for shared libraries (if it hasn't been changed and overwritten at the application level) and for shared memory (e.g. SGA).

To change "mapped base" for a Linux process, the file */proc/<pid>/mapped_base* needs to be changed where <pid> stands for the process ID. Note that this is not a system wide parameter! So in order to change "mapped base" for the Oracle database (i.e. Oracle processes), the parent shell that starts the database needs to be modified at the Linux OS level to allow it's child processes to inherit the change. The following procedure shows how this can be done.

Execute the following command to identify the process ID "pid" of the shell process used by the Oracle user that will start the database:

```
echo $$
```

As root in another shell, change "mapped base" to 0x10000000 (decimal 268435456 bytes or 256 MB) for the Oracle shell with the pid we identified above:

```
su - root
echo 268435456 > /proc/<pid>/mapped_base
```

This will tell the kernel to load shared libraries at the virtual address portion starting at 0x10000000. Now if Oracle is started with *sqlplus* in the shell used by the Oracle user for which we changed "mapped base", the Oracle processes will inherit the new base address.

Once the base address for shared memory has been changed at the Oracle level as well, more space will become available for the SGA. To accommodate the increased space for shared memory allocations by the Oracle processes, the maximum value of SHMMAX needs to be raised. This value defines the largest shared memory segment size allowed by the kernel. Since the SGA can be increased up to 2.7 GB with this method, the maximum size for SHMMAX can be

rounded to 3000000000. This will allow Oracle to allocate one large shared memory segment for the SGA. This is also what Oracle recommends.

The maximum size SHMMAX for a shared memory segment can be changed in the *proc* file system without reboot:

```
su - root
echo "3000000000" > /proc/sys/kernel/shmmax
```

Alternatively, you can use *sysctl(8)* to change it:

```
sysctl -w kernel.shmmax=3000000000
```

To make the change permanent, add or change the following line in the file */etc/sysctl.conf*. This file is used during the boot process.

```
kernel.shmmax=3000000000
```

Changing the Base Address for Shared Memory at the Oracle Level

The previous steps showed how to lower the base address "mapped base" for Oracle's shared libraries to 0x10000000 (256 MB). The following steps show how to lower the base address for shared memory (SGA) for Oracle to 0x15000000 (336 MB).

The base address for SGA (shared memory) should not be lowered to 0x10000000 at the Oracle level. As I explained in the section Address Mappings on Linux - Shared Memory and Shared Library Mapping on Linux, to prevent address range conflicts between the segments (Oracle shared libraries and Oracle shared memory), the address at which the SGA should be attached is 0x15000000. It can be lowered to 0x12000000, but this would require thorough testing. So I would not recommend it.

The following calculation shows how large the SGA can be created:

```
0xc0000000 (base address of the kernel space -> 3
GB)
- 0x15000000 (base address of SGA -> 336 MB)
-----
0xab000000 (decimal 2868903936 or 2.736 GB)
- stack space
- other memory allocations
-----
~ 2.65 to 2.70 GB
```

To lower the base address at which the SGA (shared memory) should be attached, Oracle needs to be relinked. Changing the base address for SGA can be done on Linux with *genksms*, which is an Oracle utility:

```
# shutdown Oracle
SQL> shutdown

su - oracle
cd $ORACLE_HOME/rdbms/lib

# Make a backup of the ksms.s file if it exists
[[ -f ksms.s ]] && cp ksms.s ksms.s_orig

# Modify the attach address in the ksms.s file
```

```
before
# relinking Oracle
genksms -s 0x15000000 > ksms.s
```

Rebuild the Oracle executable in the `$ORACLE_HOME/rdbms/lib` directory by entering the following commands:

```
# Create a new ksms object file
make -f ins_rdbms.mk ksms.o

# Create a new "oracle" executable
# ($ORACLE_HOME/bin/oracle):
make -f ins_rdbms.mk ioracle

# The last step will create a new Oracle kernel that
# loads the SGA at # the address specified by sgabeg
# in ksms.s:
# .set sgabeg,0X15000000
# It also backs up the old oracle executable to
# $ORACLE_HOME/bin/oracleO, # it sets the correct
# privileges for the new Oracle executable "oracle",
# and moves the new executable "oracle" into the
# $ORACLE_HOME/bin directory.
```

Rebuild the Oracle executable in the `$ORACLE_HOME/rdbms/lib` directory by entering the following commands:

```
# Create a new ksms object file
make -f ins_rdbms.mk ksms.o

# Create a new "oracle" executable
# ($ORACLE_HOME/bin/oracle):
make -f ins_rdbms.mk ioracle

# The last step will create a new Oracle kernel that
# loads the SGA at the address specified by sgabeg
# in
# ksms.s:
# .set sgabeg,0X15000000
# It also backs up the old oracle executable to
# $ORACLE_HOME/bin/oracleO,
# it sets the correct privileges for the new Oracle
# executable "oracle", and moves the new executable
# "oracle" into the $ORACLE_HOME/bin directory.
```

Now you can increase the `init.ora` parameters `db_cache_size` or `db_block_buffer` to create a larger database buffer cache. If the size of the SGA is larger than 2.65 GB, then I would test the database very thoroughly to make sure no other memory allocation problems arise.

For fun I tried to test these settings on a little test PC with 256 MB RAM and 4 GB swap space. I wanted to see if I was able to bring up a database on such a little PC. I set `db_block_buffer` to 315000 and `db_block_size` to 8192 (2580480000 bytes), and I was able to bring up a database with 2.654 GB (2850033824 bytes) SGA on this PC:

```
Total System Global Area      2850033824 bytes
Fixed Size                      450720 bytes
Variable Size                  268435456 bytes
Database Buffers               2580480000 bytes
Redo Buffers                    667648 bytes
```

Giving Oracle Users the Privilege to Change the Base Address for Oracle's Shared Libraries Without Giving them root Access

As shown above, only root can change the base address "mapped base" for shared libraries. Using

`sudo` we can give Oracle users the privilege to change "mapped base" for their own shells without giving them full root access. Here is the procedure:

```
su - root

# E.g. create a script called
"/usr/local/bin/ChangeMappedBase"
# which changes the "mapped base" for the parent
# process, the shell used by the Oracle user where
# the "sudo" program is executed (forked). Here is
# an example:

#!/bin/sh
# Lowering "mapped base" to 0x10000000
echo 268435456 > /proc/$PPID/mapped_base

# Make sure that ownership and permissions are
# correct

chown root.root /usr/local/bin/ChangeMappedBase
chmod 755 /usr/local/bin/ChangeMappedBase

# Allow the Oracle user to execute
#/usr/local/bin/ChangeMappedBase via sudo

echo "oracle ALL=/usr/local/bin/ChangeMappedBase" >>
/etc/sudoers.
```

Now the Oracle user can run `/usr/local/bin/ChangeMappedBase` to change "mapped base" for it's own shell:

```
$ su - oracle
$ cat /proc/$$/mapped_base; echo 1073741824
$ sudo /usr/local/bin/ChangeMappedBase
Password: # type in the password for the Oracle user
account
$ cat /proc/$$/mapped_base; echo
268435456
$
```

When `/usr/local/bin/ChangeMappedBase` is executed the first time after an Oracle login, `sudo` will ask for a password. The password that needs to be entered is the password of the Oracle user account.

Changing the Base Address for Oracle's Shared Libraries Automatically During an Oracle Login

The procedure in the previous section asks for a password each time `/usr/local/bin/ChangeMappedBase` is executed the first time after an Oracle login. To have "mapped base" changed automatically during an Oracle login without a password, the following can be done:

Edit the `/etc/sudoers` file with `visudo`:

```
su - root
visudo
```

Change the entry in `/etc/sudoers` from:

```
oracle ALL=/usr/local/bin/ChangeMappedBase
```

to read:

```
oracle ALL=NOPASSWD:
/usr/local/bin/ChangeMappedBase
```

Make sure `bash` executes `/usr/local/bin/ChangeMappedBase` during the login

process. You can use e.g. `~oracle/.bash_profile`:

```
su - oracle
echo "sudo /usr/local/bin/ChangeMappedBase" >>
 ~/.bash_profile
```

The next time you login to Oracle, the base address for shared libraries will be set automatically.

```
ssh oracle@localhost
oracle@localhost's password:
Last login: Sun Apr 6 13:59:22 2003 from localhost
$ cat /proc/$$/mapped_base; echo
268435456
$
```

Important Notes

When the base address "mapped base" for Oracle's processes has changed, then every Linux shell that spawns Oracle processes (e.g. listener) must have the same "mapped base" as well. This means that even shells that are used to connect locally to the database need to have the same "mapped base". For example, if you run `sqlplus` to connect to the local database, then you will get the following error message if "mapped base" of this shell is not the same as for the Oracle processes:

```
SQL> connect scott/tiger
ERROR: ORA-01034: ORACLE not available
ORA-27102: out of memory
Linux Error: 12: Cannot allocate memory
Additional information: 1
Additional information: 491524
SQL>
```

USING LARGE MEMORY PAGES (BIGPAGES)

This feature is very useful for large SGA sizes. In the following example I will show how to use and configure Linux bigpage memory area for System V shared memory segments. System V shared memory segments are allocated for SGA if "shmfs" is not used or configured for SGA.

A separate Linux memory area can be allocated to use 4 MB memory pages rather than the normal 4 kB pages. Large memory pages "bigpages" are locked in memory and do not get swapped out. This means that a whole separate bigpage memory area can be allocated for the entire SGA not to get swapped out of memory.

This means that it is very important that the bigpage memory area is only as large as needed for SGA because unused memory in the bigpage pool won't be available for other use than for shared memory allocations, even if the Linux system starts swapping. It is also important to be aware that if bigpages is set to a high value, then the available memory for user connection will be low.

Using bigpages also increases TLB (<<http://www.tc.cornell.edu/Services/Edu/Topics/Per>

[formance/SingleProcPerf/tlb.html](http://www.tc.cornell.edu/Services/Edu/Topics/Performance/SingleProcPerf/tlb.html)> "Translation Lookaside Buffers" cache hits which makes the CPUs to run more efficiently in particular with large memory configurations.

Sizing Bigpages

Oracle says that the maximum value of Bigpages should be:

Maximum value of Bigpages = $\text{HighTotal} / 1024 * 0.8 \text{ MB}$

The bigpage memory area is only available for shared memory. So if bigpages is set to a high value, then the available memory for user connection will be low. If the memory consumption for the maximum number of user connections is known, then Oracle says that bigpages can be calculated as follows:

Maximum value of Bigpages = $(\text{HighTotal} - \text{Memory required by maximum user connections in KB}) / 1024 * 0.8 \text{ MB}$

According to Oracle's white paper <http://otn.oracle.com/tech/linux/pdf/1_linuxVM_v2_accepted.pdf> "Linux Virtual Memory in Red Hat Advanced Server 2.1 and Oracle's Memory Usage Characteristics", the assumption is that 20% of memory is reserved for kernel bookkeeping.

The value for "HighTotal" can be obtained with the following command:

```
grep HighTotal /proc/meminfo
```

According to <http://www.redhat.com/advice/tips/meminfo.html>, highmem is all memory above (approx) 860MB of physical RAM. This means that "HighTotal" is the total amount of memory in the high memory region. It should now be clear that large memory pages should only be configured if enough physical RAM is available. For instance, if the server has only 512 MB RAM, then "HighTotal" will be 0 kB. And on my 1 GB RAM desktop PC, "HighTotal" shows 130992 kB.

Here are a few examples for bigpage sizes taken from <http://otn.oracle.com/tech/linux/pdf/installtips_final.pdf> "Tips and Techniques: Install and Configure Oracle9i on Red Hat Linux Advanced Server":

```
2 GB SGA 2100 MB bigpages
4 GB SGA 4100 MB bigpages
```

Using the bigpages feature for System V shared memory, the maximum size of SGA can be 5.4 GB on a machine with 8 GB RAM. Remember that HighTotal is about 7.1 GB on a 8 GB machine. If the shared memory filesystem "shmfs" is used, then the maximum size of the SGA can be increased up to 7 GB on a 8 GB machine. I'm not covering the shared memory filesystem in this article.

Configuring Bigpages

The kernel needs to be told to use the bigpages pool for shared memory allocations. The bigpages feature can be enabled for System V shared memory in the proc file system without reboot with the following command:

```
su - root
echo "1" > /proc/sys/kernel/shm-use-bigpages
```

Alternatively, you can use *sysctl(8)* to change it:

```
sysctl -w kernel.shm-use-bigpages=1
```

To make the change permanent, add the following line to the file */etc/sysctl.conf*. This file is used during the boot process.

```
echo "kernel.shm-use-bigpages=1" >> /etc/sysctl.conf
```

Setting *kernel.shm-use-bigpages=2* will enable bigpages for "shmfs" which I'm not covering in this article. Setting *kernel.shm-use-bigpages=0* will disable the bigpages feature.

The kernel needs to be told how large the bigpage pool should be. If you use GRUB, add the "bigpages" parameter in the */etc/grub.conf* file and set the maximum value of bigpages as follows. In this example I will set *bigpages* to 2100 MB for the SMP kernel 2.4.9-e.25 that is started on my database server:

```
default=1
timeout=10
splashimage=(hd0,1)/boot/grub/splash.xpm.gz
title Red Hat Linux (2.4.9-e.25enterprise)
    root (hd0,1)
    kernel /boot/vmlinuz-2.4.9-e.25enterprise ro
root=/dev/hda2 hdc=ide-scsi
    initrd /boot/initrd-2.4.9-e.25enterprise.img
title Red Hat Linux Advanced Server (2.4.9-e.25smp)
    root (hd0,1)
    kernel /boot/vmlinuz-2.4.9-e.25smp ro
root=/dev/hda2 hdc=ide-scsi
    bigpages=2100MB
    initrd /boot/initrd-2.4.9-e.25smp.img
title Red Hat Linux Advanced Server-up (2.4.9-e.25)
    root (hd0,1)
    kernel /boot/vmlinuz-2.4.9-e.25 ro
root=/dev/hda2 hdc=ide-scsi
    initrd /boot/initrd-2.4.9-e.25.img
```

After this change the system needs to be rebooted:

```
su - root shutdown -r now
```

After a system reboot, the "MemFree" value (free system memory) in the */proc/meminfo* is subtracted by 2100 MB in this example. The 2100 MB show now up in the "BigPagesFree" which means that 2100 MB are now in a separate allocation area:

```
grep MemTotal /proc/meminfo
grep BigPagesFree /proc/meminfo
```

Note that if you configure "bigpages" in the */etc/grub.conf* file and reboot the system that "BigPagesFree" in */proc/meminfo* will be 0 KB if "HighTotal" in */proc/meminfo* is 0 KB and if */proc/sys/kernel/shm-use-bigpages* is set to "1".

MAKING OTHER PERFORMANCE RELATED CHANGES

Disabling Unneeded Background Processes

X should not run unless you need to. You can stop X by switching to runlevel 3 with the following command:

```
init 3
```

To switch back to runlevel 5 so that X comes up again, run:

```
init 5
```

To set the default runlevel permanently to 3 so that X doesn't come up with the next reboot, change the following line in */etc/inittab*

```
id:5:initdefault:
```

so that it reads:

```
id:3:initdefault:
```

You can check for other unneeded background processes by running the command:

```
/sbin/chkconfig -list
```

To temporarily disable e.g. *ybind*, run:

```
su - root
service ybind stop
```

To permanently disable *ybind*, run:

```
chkconfig ybind off
```

ORACLE ERRORS AND PROBLEMS

The intention of this section is to describe errors and problems that can occur in connection with the changes covered in this article.

For errors regarding the installation of Oracle software and regarding the creation of a database, see <http://www.puschitz.com/InstallingOracle9i.shtml> "Oracle Installation Errors"

ORA-3133 errors and attach errors

Cause(s):

- Running an Oracle binary that has a lower SGA base, but */proc/proc/<pid>/maps* has not been adjusted as well.
- SHMMAX value has not been increased large enough.

```
SQL> startup
ORA-03113: end-of-file on communication channel
SQL>
```

Cause(s):

- A too large SGA has been configured

- SHMMAX value has not been increased large enough.
- Oracle has been relinked with a lower SGA base address but "mapped base" has not been lowered for the shell at the Linux OS level.

```
SQL> startup
ORA-27102: out of memory
Linux Error: 12: Cannot allocate memory
Additional information: 1
Additional information: 262148 SQL>
```

- Cause(s):
- This error message comes up if the SGA size is too large.

ORA-01041: internal error. hostdef extension doesn't exist

- Cause(s):
- If this error comes up and the database is not up, then remove all shared memory segments from the Linux OS.

USEFUL LINUX PERFORMANCE UTILITIES

top Utility

This utility shows CPU consumption, memory consumption, and "top" sessions on the Linux server:

```
top
```

Load Averages:

The first line of the *top* output shows you a series of three "load average" numbers. These numbers describe the load on the system. The load average is the average number of processes that are waiting in the queue for CPU time (including processes that are waiting for I/O) for the past 1, 5 and 15 minutes.

For example, if you run 3 non-interactive processes that are not waiting for input or I/O, then you can expect the average load to be 3. To illustrate that, run the following command in 3 different shells on a server that is not being used:

```
while [ 1 ]; do str="x"; done
```

This loop will use up all the CPU time that it can get since it's not doing any I/O and it is not waiting for input. Now wait for about 2-3 minutes and you will see that the average load for the last 1 minute will increase to be 3 and higher. It will be a little bit higher than 3 since there are other processes running on the system.

In general, a number less than 1 is ideal. A load average value of 3 is high. And a value of 10 is definitely a heavily loaded system where you can expect delays.

You can also use the *load* command to display real-time text mode graph on the "load average".

CPU States:

It shows the load on each processor - the percentage of CPU time in user mode, system mode, niced tasks, and idle.

The "user" percentage shows how much processing time the CPU is spending on user processes, and the "system" percentage shows how much processing time the CPU is spending in the system (kernel). Niced tasks are only those processes whose nice value is negative. And note that the processing time for niced processes will also be counted in system and user time, so the total will be more than 100%.

However, the best indicators of a stressed CPU is the load average which I described above.

Sessions:

This section shows the top sessions (Linux processes) in terms of CPU utilization.

sar Utility

sar stands for System Activity Reporter.

CPU Usage:

To check CPU usage over time, run:

```
sar -u
```

This command is useful if you want to see overall CPU consumption over time.

%user shows the percentage of CPU utilization at the user level (application).

%system shows the percentage of CPU utilization at the system level (kernel).

To check CPU usage 10 times with a time interval of 3 seconds, run:

```
sar -u 3 10
```

Swap Activity:

To check swap activity over time, run:

```
sar -W
```

This command is useful if you suspect memory shortages.

pswpin/s shows the total number of swap pages the system brought in per second.

pswpout/s shows the total number of swap pages the system brought out per second.

These numbers should be low. If not, you need more RAM.

To check swap activity 10 times with a time interval of 3 seconds, run:

```
sar -W 3 10
```

I/O Activity:

To check physical disk I/O activity over time, run:

```
sar -b
```

This command is useful if you suspect that the database is I/O bound.

See manual pages for more information.

To check I/O activity 10 times with a time interval of 3 seconds, run:

```
sar -b 3 10
```

vmstat Utility

This utility provides a report that covers process activity, paging, memory usage, disk I/O, and CPU usage.

To create 5 reports with a time interval of 3 seconds, run:

```
$ vmstat 3 5
procs memory swap io system cpu r b w swpd free buff cache
si so bi bo in cs us sy id
0 0 0 186460 7416 9424 45272 1 4 25 35 126 33 3 0 96
0 0 0 186460 7416 9432 45272 0 0 0 17 103 18 0 0 100
0 0 0 186460 7288 9440 45272 0 0 0 73 104 23 4 1 95
0 0 1 186460 7288 9440 45272 0 0 0 5 102 12 0 0 100
0 0 0 186460 7288 9440 45272 0 0 0 8 102 14 0 0 100
```

See man pages for more information.

ORACLE LINUX MANAGEMENT

Determining Which Semaphore Sets and Shared Memory Segments Belong to Each Oracle Database or Instance

When Oracle hangs or crashed or when Oracle was killed, then sometimes you will see that shared memory segments and/or semaphore sets have not been released or removed by the Oracle background processes. It is important to make sure that the semaphore sets and shared memory segments are released at the Linux OS level before the database or instance is restarted.

Running *ipcs* will only show you which semaphore sets and which memory segments are owned by the Oracle user account. If you have only one database running on your server, then you can simply use the IDs of all shared memory segments and semaphore sets that belong to the Oracle user account and release them via *ipcrm*:

```
su - oracle
$ ipcs
----- Shared Memory Segments -----
key shmid owner perms bytes nattch status
0x00000000 0 root 600 196608 2
0x00000001 32769 root 600 655360 2
0x00000000 458755 oracle 660 4194304 0
0x00000000 491524 oracle 660 33554432 0
```

```
0x00000000 524293 oracle 660 33554432 0
0x00000000 557062 oracle 660 33554432 0
0x00000000 589831 oracle 660 33554432 0
0x00000000 622600 oracle 660 33554432 0
0x00000000 655369 oracle 660 33554432 0
0x00000000 688138 oracle 660 33554432 0
0x3ecee0b0 720907 oracle 660 4194304 0

----- Semaphore Arrays -----
key semid owner perms nsems status

----- Message Queues -----
key msgid owner perms used-bytes messages
$
```

To release all shared memory segments that are owned by the Oracle user as listed above, run:

```
$ ipcrm shm 458755 491524 524293 557062 589831
622600 655369 688138 720907
```

The command for releasing semaphore sets is:

```
ipcrm sem <semid>
```

But if you have more than one database or instance running on the Linux servers, then *ipcs* will NOT show you the semaphore sets and shared memory segments that are owned by each database or instance. The following steps can be used to find the right IDs for each database or instance:

```
$ su - oracle
$ sqlplus /nolog

SQL> oradebug setmypid
Statement processed.
SQL> oradebug ipc
Information written to trace file.
SQL> select value from v$parameter where name =
'user_dump_dest';
VALUE
-----
/opt/oracle/admin/test/udump

SQL>
```

On my test server, the *oradebug ipc* command created a file called *test_ora_6626.trc* in the *USER_DUMP_DEST* directory */opt/oracle/admin/test/udump*. The name of the created trace file is *\$ORACLE_SID_ora_<pid>.trc* where *<pid>* stands for the process ID of the Oracle foreground process in a non-MTS environment that's talking to *sqlplus* here. If you are not sure about the name of the file that was created, run *ls -lrt* to see the timestamp of the latest trace file created in the *USER_DUMP_DEST* directory.

When you open the trace file (in my example *test_ora_6626.trc*), you can find the semaphore ID for this database after the line "Semaphore List=". Here are the semaphore sets on my test box for the Oracle database:

```
/opt/oracle/admin/test/udump/test_ora_6626.trc:
[SKIP]
Maximum processes: = 150
Number of semaphores per set: = 154
Semaphores key overhead per set: = 4
User Semaphores per set: = 150
```

```

Number of semaphore sets:      = 1
Semaphore identifiers:         = 1
Semaphore List=
98304
----- system semaphore information -----
Shared Memory Segments -----
[SKIP]

```

To release all semaphore sets that are owned by the database as listed above, run:

```
ipcrm sem 98304
```

And here are the shared memory IDs on my test box for the Oracle database:

```

/opt/oracle/admin/test/udump/test_ora_6626.trc: [SKIP]
Area #0 'Fixed Size' containing Subareas 0-0
Total size 000000000006e078 Minimum Subarea size 00000000
Area Subarea Shmid Stable Addr Actual Addr
0 0 1671186 0x000000050000000 0x000000050000000
Subarea size Segment size
000000000006f000 0000000000400000
Area #1 'Variable Size' containing Subareas 1-7
Total size 000000000e000000 Minimum Subarea size 01000000
Area Subarea Shmid Stable Addr Actual Addr
1 1 1703955 0x000000051000000 0x000000051000000
Subarea size Segment size
0000000002000000 0000000002000000
Area Subarea Shmid Stable Addr Actual Addr
1 2 1736724 0x000000053000000 0x000000053000000
Subarea size Segment size 0000000002000000 0000000002000000
Area Subarea Shmid Stable Addr Actual Addr
1 3 1769493 0x000000055000000 0x000000055000000
Subarea size Segment size
0000000002000000 0000000002000000
Area Subarea Shmid Stable Addr Actual Addr
1 4 1802262 0x000000057000000 0x000000057000000
Subarea size Segment size
0000000002000000 0000000002000000
Area Subarea Shmid Stable Addr Actual Addr
1 5 1835031 0x000000059000000 0x000000059000000
Subarea size Segment size 0000000002000000 0000000002000000
Area Subarea Shmid Stable Addr Actual Addr
1 6 1867800 0x00000005b000000 0x00000005b000000
Subarea size Segment size
0000000002000000 0000000002000000
Area Subarea Shmid Stable Addr Actual Addr
1 7 1900569 0x00000005d000000 0x00000005d000000
Subarea size Segment size
0000000002000000 0000000002000000
Area #2 'Redo Buffers' containing Subareas 8-8
Total size 0000000000a3000 Minimum Subarea size 00000000
Area Subarea Shmid Stable Addr Actual Addr
2 8 1933338 0x00000005f000000 0x00000005f000000
Subarea size Segment size
0000000000a3000 0000000000400000
Area #3 'skgm overhead' containing Subareas 9-9
Total size 000000000001000 Minimum Subarea size 00000000
Area Subarea Shmid Stable Addr Actual Addr
3 9 1933338 0x00000005f0a3000 0x00000005f0a3000
Subarea size Segment size
0000000000001000 0000000000400000
[SKIP]

```

To release all shared memory segments that are owned by the database as listed above, run:

```
ipcrm shm 1671186 1703955 1736724 1769493 1802262
1835031 1867800 1900569 1933338
```

To verify if the shared memory segments and semaphore sets have been released, run:

```
ipcs
```

HARDWARE RECOMMENDATIONS

It really depends on what kind of database you want to setup and run, how large the database is etc.

But people keep asking me what I would recommend. If you want to get a feeling how well Oracle9i (non-RAC system) runs on Linux/Intel systems, and if you don't want to spend "too much money", here is what I would buy:

- 2-way server, 2.4GHz Xeon
- 4 GB RAM; RAM is cheap and gives you usually the biggest "bang for the buck".
- Large Internal Ultra SCSI disks with a hardware RAID controller card.

REFERENCES

<<http://otn.oracle.com/tech/linux/content.html>>
Oracle's Linux Center

<<http://www.redhat.com/whitepapers/rhel/AdvServerRASMPdfRev2.pdf>> An Overview of Red Hat Advanced Server V2.1 Reliability, Availability, Scalability, and Manageability (RASM) Features

<http://otn.oracle.com/tech/linux/pdf/1_linuxVM_v2_accepted.pdf> Linux Virtual Memory in Red Hat Advanced Server 2.1 and Oracle's Memory Usage Characteristics

<http://otn.oracle.com/tech/linux/pdf/installtips_final.pdf> Tips and Techniques: Install and Configure Oracle9i on Red Hat Linux Advanced Server

<<http://otn.oracle.com/tech/linux/pdf/9iR2-on-Linux-Tech-WP-Final.PDF>> Oracle9iR2 on Linux: Performance, Reliability and Manageability Enhancements on Red Hat Linux Advanced Server 2.1

Delivering Leading TPC-C Figures with Red Hat Linux Advanced Server (Red Hat Webcast Tuesday, 22nd October, 2002)

<<http://www.bookpool.com/.x/4ycq2gjumn/sm/0596002130>> Understanding the Linux Kernel, 2nd edition

<<http://www.bookpool.com/.x/4ycq2gv078/ss/1?qs=oracle8i+internal>> Oracle8i Internal Services for Waits, Latches, Locks, and Memory

Copyright Notice: This article may not be published, sold, reproduced or copied in whole or in part without obtaining permission first. But you are welcome to put links from your site to the article.

The information provided in this article shows how I tuned Linux/Oracle on my server(s) and is distributed AS IS. Every effort has been made to provide the information as accurate as possible, but no warranty or fitness is implied. The use of this information described herein is your responsibility, and to use it in your own environments do so at your own risk.

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

[URL:http://www.puschitz.com/TuningLinuxForOracle.shtml](http://www.puschitz.com/TuningLinuxForOracle.shtml)

AUUGN Book Review

Reviewed by Greg Lehey <greg.lehey@auug.org.au>

JEFF DUNTEMANN'S DRIVE-BY WI-FI GUIDE PARAGLYPH PRESS, 2003

The cover of Jeff Duntemann's Drive-by Wi-Fi Guide promises:

- ◆ Build your own low-cost Wi-Fi network.
- ◆ Understand the risks and master Wi-Fi security.
- ◆ Learn the secrets of gain antennas, bridging, and wardriving.

One thing it doesn't say is that the book is about wireless networking; you need to read the rear cover carefully to discover that. If, like me, you're more comfortable with standard numbers like 802.11, you might overlook the book altogether.

The somewhat over 400 pages of large, well-spaced text are divided into four sections: network basics, wireless hardware, security and wireless networks as a hobby. This last section includes information on building antennas out of tin cans and "wardriving", driving around looking for wireless networks. It was published earlier this year, and it includes information about 802.11a, the 54 Mb/s standard for the 5 GHz band. It can be forgiven for not knowing about 802.11g, the 54 Mb/s standard for the 2.4GHz band which rather suddenly popped up in the middle of the year.

The book starts off by describing networking in general and wireless networking in particular, unfortunately very inaccurately. It describes the Internet Protocols, but without explaining the differences between link and transport, between bridging and routing. The latter issue is of particular interest for wireless networking, and it was the main reason I wanted to read a book on wireless networking: does the wired network see your MAC address (bridging), or does it just see IP datagrams (routing)?

The section "Designing your WiFi Network" looked like another useful topic, particularly since I wanted to deploy more than one access point, which requires an ESS (Extended Service Set). Unfortunately, it doesn't say anything worthwhile about the topic, and the section finishes with a suggestion that it's not the end of the world to have to "log in" again as you move from one place to another. Tell that to somebody with six NFS mounts.

The author goes to some trouble to explain the constraints on the network topology, apparently because he doesn't know (or at least want to talk) about ARP. In all probability there's more to the issue than I understand--that's why I read the book--but the text didn't convince me.

The next part of the book covers "current"

hardware. This seems of limited utility: only a couple of months after publication, much of the hardware is already obsolete. This section does have some useful information about the physical layout of network equipment, however.

Much of the book seems to address latter day radio amateurs: it goes into detail about setting up long-distance wireless links, building antennas and locating foreign wireless networks. That's certainly an aspect of mild interest, despite certain concerns about the legality of the advocated actions, but you only need to look at the offerings in a large computer supermarket such as Fry's to realise that this is not the area of wireless networking which interests the general public.

Interspersed between details of finding public "hotspots" and "The Tin Can Bandwidth Expander Project" are two separate sections on security, apparently derived from different sources. The first is an overview of security issues, including a page or two of good information on the kinds of security exploit you might see on any network, and also some information on WEP. The second part describes how to navigate the security menus on a Linksys "residential gateway", the driver for an (obsolete) Orinoco Gold Card under Microsoft "Windows" XP. It's not clear how useful this information is to people with different hardware or software.

The book does its name every justice: it's the opinion of one person (Jeff Duntemann), it's superficial ("drive-by") and it's full of buzzwords ("WiFi"). It's far too inaccurate and superficial to be useful to a technical audience. The book I was looking for would have had the title "Practical 802.11 wireless networking". It didn't answer my question: "How do I set up an EBSS?" to my satisfaction: I was left with the conviction that the answer it provided was wrong, which was confirmed by subsequent experience. What about the home user? They might be satisfied with it, but replacing facts with buzzwords doesn't help anybody.

Play Encoded DVDs in Xine

Author: LeaAnne Kolp

First of all, you'll need to download the plugins:

- xine_d4d_plugin-0.3.2.tar.gz
- xine-d5d-0.2.7.tgz
- xine-lib-0.9.12.tar
- xine-ui-0.9.12.tar

These plugins will ONLY work with the xine-lib and xine-ui-0.9.12. If you get 0.9.13 it will NOT work.

After you download those, switch to root

```
[tux@linux tux]$ su
Password: *****
```

Then you'll have to move the files that you just downloaded to your /root/ directory. Do this by typing in the following at the command prompt.

```
mv *.tar.gz /root/
```

If that doesn't work, then just type out the following:

```
mv xine-lib-0.9.12.tar.gz /root/
```

Do this for each of the files. After you get that done, then switch to your /root/ directory by typing in the following:

```
cd /root/ type: ls
```

And you'll get a listing of all the files in your /root/ directory. Now for the good part.

Now to gunzip and untar it

To do this, type in the following:

```
gunzip -d xine-lib-0.9.12.tar.gz
```

Switch to that directory by typing the following:

```
cd xine-lib-0.9.12
```

Now type in:

```
ls
```

Now that you're in the directory, you'll have a README file and INSTALL file. ALWAYS read the README file. No matter how many times you've done this before, something might have changed. If the README doesn't tell you anything read the INSTALL file.

To do this, type in:

```
more README
```

(just like it is in the directory, if you don't type it identical, it won't access it)

```
more INSTALL
```

Normally, a typical installation is done by typing in these commands:

```
./configure  
make  
make check  
make install
```

Again, always read the README. Each distribution of Linux is different and therefore the installation instructions could be different.

Keep repeating the above steps until all 4 files are installed.

Then type in:

```
tar -xvf xine-lib-0.9.12.tar
```

Now type in:

```
ls
```

You'll see a directory (in dark blue) with the name:

```
xine-lib-0.9.12
```

Then type in the following to update your drive:

```
updatedb
```

That could take awhile depending on your drive. When that's done, you'll have to locate xine. To do so, type in the following:

```
locate xine
```

It usually puts it in /usr/local/bin/ but to be on the safe side, locate it. :)

Once you have it located, until you add it to your menus, type in the directory of where it was. So if it was in /usr/local/bin/xine you would type in: /usr/local/bin/xine

That would start the program running if that's where it was located.

Now here's the tricky part that you'll have to play with and figure out on your own. When xine comes up, you'll see the d4d and d5d buttons at the bottom. When you put a dvd into the dvdrom drive you'll have to click on either the d4d or d5d button to get it to play the encoded dvd.

Unfortunately, I don't know which one will work with the dvd you put in.

Some dvds take the d5d, others take the d4d, you'll just have to play around with it and experiment to find the one that's right.

What I've started doing is when I put a dvd in and find out which plugin works (i.e. d4d, d5d) I write it down, so I know and I don't have to play games with it to figure it out! :)

Congratulations! You've just gotten the plugins to work and now you can sit back and enjoy the movie!

Hi, my name is LeaAnne and I've been Windows Free since March 2003.

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

<http://www.tldp.org/LDP/LG/current/kolp.html>

Interview: Australian Sun HPC Team

Interviewed for AUUGN by: Con Zymaris <auugn@auug.org.au>
Transcribed by: Sahand Shoghi <sahand@cybersource.com.au>

Cast of Characters

Richard Smith – HPC Specialist in Australia

Con – AUUGN Editor

Steve Davies - Lead Systems engineer on the project

Justin Glen - Senior Industry Account Manager

Con: Ok so firstly let me thank you for taking time to join with me this morning.

Richard Sure!

Con: We're discussing the large super-cluster deployment you recently did for the University of Queensland. What can you tell me generally about the deployment? From just an overall kind of perspective. Is this something relatively a new kind of line of work for Sun? I understand obviously it's a cluster kind of arrangement this particular solution, wasn't it?

Richard Yes it is!

:

Con: There are some things that obviously Sun is fairly well known in like the mainframe Unix type stuff but for doing this HPC type stuff, is that fairly new for Sun in Australia?

Richard Sun's been in the HPC space here for some considerable time and this, this particular area of HPC has to do with what we call parallel computing

Con: Yep!

Richard Now the nature of this is that, there are certain programs, or codes as their called in HPC that lend themselves to parallel execution, now what this means is that instead of working on a program from end to end in a serial fashion, it is possible to dispatch the program against the machine and have the lot operated on in parallel and the end of that execution for all this information to come back together to arrive at the outputs at the other end.

Con: Yep, indeed

Richard Now this is what we mean by parallel computing and this implementation at the University is specifically attuned to that type of computing.

What the sort of codes and the sort of research the University does, will lend itself to this parallel/parallel execution, if you like.

Con: So, in terms of their problem spaces are very highly parameterized, you can actually sort of break it down to discreet chunks.

Richard Yes it is, but in fact the machine does a lot of that for you, in terms of your submitted code or problem and it will dispatch it across the multiple processors for execution.

Con: Fair enough, ok, so that obviously leads into the next question. Why was that particular hardware architecture (the cluster approach) selected in lieu of any alternative types of architecture?

Justin: Well uh it's Justin here I could respond to that; The University has established some investment in other types of architectures already, they have some large SMP machines, and they were looking for a cluster, a low-wide cluster which would suite a particular type of application that many of the engineers are running and also used by the chemical people. And so this complemented their existing infrastructure. And obviously the University of Queensland are also part of the APAC (Australian Partnership for Advanced Computing, <http://www.apac.edu.au/>) relationship which has a centralized facility based in Canberra, so they have a low-wide cluster, that's a cluster of four CPU nodes or eight CPU nodes, so they felt that this cluster of two CPU nodes complemented all of their existing infrastructure, so it fits very nicely.

Con: Fair enough, obviously if you draw up your standard sort of Gartner quadrant where you know you get particular sweet spots for for hardware and obviously two nodes are a lot cheaper to deploy then four and definitely eight-way type nodes. I presume it's an x86 type of arrangement?

Justin: Yes, that's right. So this is a 128 way, 2-CPU, 2.8 Ghz Xeon arrangement. As you increase the number of CPUs per node the cost per CPU increases and so conversely as you reduce the number of CPUs in each node the total cost the CPU reduces, and you can take that to the extreme and have a large number of single CPU machines. It turns out that the sweet spot for us is a two CPU machine and that gives effectively in our product range, that would give, if you wanna use theoretical peak teraflops, which is what we measured it initially at, if you do the mathematics that is 256 CPUs times 2.8 Ghz times 4 floats per instruction it works out to be 1.4 teraflops in performance.

Con: Which I presume just based on that would easily get this box into the top 50 or top 100 in *top500.org*?

Justin: Well that's a moving target issue now.....

Con: Of course

Justin: And the last target was published in June or July, I forget it now, and this system would appear to make the top 100 on that list.

Con: Right

Justin: When we completed the benchmark, the *Linpack* benchmark on this system and to submit unto the new list in November then we'd expect it to be fairly close to the top 100 in the world and if you look at the theoretical peak then it would be number one in Australia at the moment from what we can tell.

Con: Right

Justin: There maybe some systems out there that we don't know about and I think if you look at actual performance on the Linpack *top500.org* number it would be fairly close to number one in Australia if not number one . It really depends on how the benchmark runs.

Con: Understood, and when the was at the time the snapshot was taken

Justin: Yeah that's right and the current plan is to have that submission done soon. I don't know what the deadline is at the moment but it is usually a month before the list comes out, the list normally comes out early November so we would like to get the benchmark completed by early October.

Con: Fair enough. To make that cut off.

Justin: Yeah!

Con: Of course. Now in terms of APAC, which you mentioned beforehand, is really is moving along. There are various components of APAC in the various states across Australia in the University and Government sector I know down here VPAC have just rolled out a system, and I think there's been a couple of systems in Queensland over the past twelve months, just recently the University of Adelaide has had another fairly major system roll out.

If you go back some years in the US, to the National Center for Supercomputing Applications (NCSA), which had a fairly dominant role in gluing together these otherwise disparate researchers and research groups and trying to conjoin them synergistically into using the same platforms and having them leverage off each others work, How much of that currently happens in Australia?

Is APAC operating just at a bearucratic level? Are we looking at a situation where, X many compute cycles of time in Queensland are available for reuse to South Australia? Is what actually happens more at that kind of technical networking level rather than the people level?

Justin: Well there is a political answer to that and there is a technical answer and let me respond to the technical question.

The system that's going into U.Q. is capable technically of participating in a in a national or a world-wide grid for that matter, and the trick and the secret really is to join that grid into the national system and to have those grid products, the software products that are driving them either the same or talking to one and other. The system that's going to U.Q. that we're putting there has a product which we call Grid Engine Enterprise Edition and that came from a purchase that Sun Microsystems did a couple of years ago where we bought the company that made the product and now we distribute that product free of charge to universities, that's the base product, the Enterprise Edition we charge for and we have the commercial range with universities of Queensland where they get access to that product.

So the Grid Engine is a software product that allows a user to dispatch their jobs to any or all of the 256 CPUs that are in that specific computer and it could extend to beyond that facility, if it was linked to other facilities.

Of course there is a political response which is that the university has paid for the cycles and this supercomputer, and none of the other universities I believe would have contributed to that facility so there is probably no right for them to get access to these compute cycles.

Con: Obviously, But your looking at a situation like for example much like the national Grid for electricity you can have State X offering for sale to State Y you know X amount of mega wattage if you like, so that is obviously the end-point for this kind of Grid style computing and there is obviously an increasing amount of interest nationally. I believe the National Office of The Information Economy (NOIE) are seriously looking at assisting in the development of this industry; the computing grid industry and that's something I don't know if you guys have already been briefed about or asked to participate in?

Justin: Yeah well that's not something where I'm looking at at the moment, I can say that as a company Sun is more focused on making HPC more usable and that is allowing users not just IT or AUUG members to use the facility, but it's to allow any researchers simple access to a system and a way that we're approaching this problem, well one approach would be to take your academics and put them in a two or three-day Unix class and give them some Unix skills which is a possibility, in fact University of Queensland are doing that right now for other systems.

Our approach is somewhat different, that is to tackle the problem using some smart technology and we have a product which is called Technical Computing Portal (also known as HPC portal.) And Sun is dominant in the portal and the intranet market place, so what we want to do is link our expertise in that area to our expertise in HPC and we've developed a product and a suite of tools which will assist the I.T. people at the University of Queensland. They've signed up to do this for this cluster, to make it available through a portal, so it's simply a matter of a user going to a particular web address and they click on the applications that they wanna run and there will already be a number of pre-loaded applications installed on the system. So they click on the application then they click on the data input which could be a file on their PC, on their desktop so they can browse their own C: drive and nominate a file to be used as the input or they could indeed nominate a file that's on the network at University of Queensland or somewhere other than their own University and then they would click on which parameter options that they would want to use at runtime, select the number of CPUs that they would like to run this application on, then submit the job.

And then the portal takes care of all of the hard bits, that is the authentication the authorization, it looks at moving the files, getting the data into the right place and it also distributes the output data to the right people. So that way someone can use this system without having Unix expertise .

Con: Yep, that sounds like a good idea, Now on this functionality specifically; obviously, at some point, somewhere either Sun developer or a UQ researcher would need to actually develop and test, deploy the appropriate plugin middle-ware which sits behind or in the portal, if you like, that actually does the data munging, obviously through calling existing infrastructure to pull this data into the Grid engine. So who would develop that, or is there a set of pre-existing compute requirements that you have actually specified and they've already been developed?

Steve: Ok Con yeah it's Steve here. What you've described is pretty well how it happened, there is the two pieces; There is the front end which Justin has described which is the portal, give you like a web based interface to it, it's a bit like doing web based email...

Con: Yep

Steve: Ok so you've got your browser there and you can submit these pre-defined jobs.

The other piece of the software is the other item that Justin mentioned and that's the Sun Grid Engine Enterprise Edition and this is what does the actual processing and distribution of the code. It's a piece of software that looks after distributing the code across the array or across the cluster if you like and then compiling the results back together for display back at your web-browser.

Now I'm not sure whether that's where you were headed with that sort of discussion. There will of course need to be some human interaction in terms of setting up the Grid engine in terms of defining what codes will be available to the various researches.

Con: Yep, of course, for example one problem might be deploying a fairly standard simulated annealing process in code and obviously that's a fairly well known and understood problem, somebody would have developed the necessary module if you like that interfaces with the grid engine to distribute that kind of problem request. But there might be other ones where that particular university might do work for example in aerodynamics for which they've developed their own formula etc., so obviously someone at that point will need to develop the necessary module that sits behind the portal engine and in front of the grid engine to offer the appropriate parameters as a web interface; accept the user input, put this thing into some kind of compute engine mechanism if you like, which calls the appropriate back end grid engine components for doing all the necessary delivery of the particular codes the the nodes. Because obviously there is a whole world of different types of mathematics that can be done on this kind of cluster system and they can't have all been thought of before hand as a drop list option set.

Steve: Uh sure yeah, and to give an example one of the things that UQ will be looking at, is analysis of Hydrogen, I guess you'd be aware of the possibility of hydrogen becoming a viable fuel for motor vehicles.

Some of the challenges around that are to do with storing the Hydrogen, handling the Hydrogen, transferring Hydrogen from storage devices into the motor vehicles. Because it's a fairly volatile sort of substance, there are a number of issues around that, and they need to be resolved before it can become a viable fuel.

Con: Yes.

Steve: But for example some of the unis might be doing is using this super computer to analyse the molecular properties of Hydrogen and simulate the way in which Hydrogen works with other chemicals.

Con: Right

Steve: Now these are the sort of research codes that they would execute and then hopefully come up with methods of translating that into a commercial solution to handling Hydrogen.

Con: Of course.

Next question I have for you is the operating system platform for this, I presume a lot of the x86 based systems generally deploy with something like Linux I presume that's the solution with this one as well?

Steve: Yes that's right, it will be a Red Hat.....

I can't just tell you which version that is because we're supplying the system and the customer has decided to download their own release of Red Hat, so I think it would be premature for me to say what that would be because they still going through the installation and the customisation at the moment.

Con: Fair enough.

So obviously it's just a Linux 2.4 type kernel and eventually go up-to 2.6 when it comes out the next few months.

Steve: Yes

Con: Now obviously from Sun's perspective it's got broad based support for Solaris on x86 and also for Linux on x86.

Steve: That's right.

Con: Is there any specific either/or proposition to which one you would deploy for this kind of cluster based on the dual Xeons and so forth? Obviously when you start talking the high end, Sun 10000 and 15000 Star Cats, that's a different kettle of fish as far as Sun's concerned and all based on Solaris.

At the cluster level, the application infrastructure the grid engine, the portal server and so forth, they don't really care what the underlying Operating System is - Solaris, Linux it pretty much runs?

Steve: That's right, so the grid engine, when we acquired that organization, had grid engine, they had that product running on multiple platforms and we kept it the same.

Con: Understood.

Steve: So it's in-different to whether it's Solaris or....

Con: From memory, I think it's based on Java anyways?

Steve: Uh yes thats right

Con: Righto.

Steve: And that formula's is used in Java so really it's all in-different about what sort of systems are underline that.

And your question about MPI code, I think your asking about that earlier on we have a product called HPC cluster tools (<http://www.sun.com/servers/hpc/software/>) which came from our acquisition of Thinking Machines, it was a massively-parallel system which come out years ago....

Con: I remember **Grin**

Steve: One of their systems featured in the, hmmm, now which movie was that hmmm.....

Con: Lots-n-lots of blinking lights...

Steve: Yeah thats right!!! It was hmm, the dinosaur one...

Con: Jurassic Park!

Steve: I have some fond memories of that system, the hardware part of that company went elsewhere and the software part was acquired by Sun, and part of the software turned into a product that we have called HPC cluster tools which has all the MPI libraries.

Con: Right

Steve: So the customer will either use that product or use Open Source MPI libraries which sits underneath the grid engine and that would be what the academics will write their code to.

Con: The Thinking Machine firm and the Connection Machine were created by a fellow called Danny Hillis back in the mid to late 80's. He was doing a lot of work with another fellow called Steven Wolfram, primarily on 2 and 3D Cellular Automata (CA)

Steve: Yep..

Con: And Wolfram later on went on to create in the late 80's Wolfram Research, wrote *Mathematica*, has recently published this big massive tome on his research on CA

Steve: Ok

Con: It's interesting stuff. I didn't realize that Sun ended up with at least a component of the former Thinking Machine's stash

Steve: Yes, and we obviously ended with a component of Cray as well

Con: Yes, and this is part of the NUMA work

Steve: So Sun does have a long history of HPC. I suppose and It's a hidden secret in Sun.

Con: In past I remember when I first started reading about supercomputers, the one or two of these which were floating around, this is before Cray2 came, it seems to be that for many many years, the locus of development of supercomputer technologies in the world was Minneapolis. That's where Cray, Seymour Cray had been working at Control Data Corporation in the same city, that's where many of the mini super computer companies were located.

It seems to be in the past five, maybe ten years in particular that to a large extent all the former super computer companies particularly, in the US, have been swallowed up by much much larger organizations like SGI and IBM and obviously Sun. And that kind of specialized or specialist super computer fraternity has pretty much gone, the only major company I think that's still in existence, although in dramatically modified form, is remnants of Cray Research

Steve: Cray?

Con: Cray still exist as an entity but it's not really the same company that was there ten years ago when they were acquired SGI.

Whats your understanding or take on the industry side of this? It seems to be the case that obviously supercomputer's have become fairly big, complex and powerful, and there seems to be like a bit of bifurcation on the technology, on the one hand you've got (particularly) the Japanese organisations (NEC, Hitachi) pushing expensive vector array processors

Steve: Yep

Con; And on the other hand you've got mostly the American firms who are mostly pushing a clustered kind of approach whether it's a HPC arrangement you guys have which is fairly similar to the kinds of deployments that IBM and HP do and there is theslight variation on that with SGI's platforms, the Altix 3000 in particular.

There seem to be two main classes of problems to solve with HPC systems. You can very much do most things with the cluster kind of solution but there is another class of problem that is far more suited to be solved by using the vector processing engines. However, those vector processing engines become incredibly difficult to improve and to push the state of technology on. They're far more expensive but they do score runs on certain class of problems. Is that pretty much as Sun would see it?

Steve: Yeah I think you probably described it very well, there are two classes of problems, you got your single threaded end to end type of transaction where you start at the start and you need to go through a specific sequence of events to arrive at the end and for that type of problem there is not a lot of opportunity to parallelise. Whereas there other sorts of problems that can be split up rather easily and you can have computations going on in parallel simultaneously if you like, that eventually arrive at an end-point after they have been happening separately, so it does in the end I think come down to the fact that there are two sorts of problems and the relative mirror of each approach and use that approach that best fits that problem.

Con: Of course yeah.

Justin: Justin here. My take on this as that Sun has been successful in the HPC arena promoting common hardware that is more general purpose hardware (read lower costs) and more economic for HPC use. As opposed to building specialized hardware which is designed for HPC.

That has some advantages because that means that our company can be commercially viable because we're using the same hardware for banks, telecommunications organisations as are we are for HPC. The way we can tailor it by putting software applications on there such as HPC cluster tools, grid engine, technical compute portal which make the technology more usable to most users and that's been the strategy that Sun's had and we're number two or number three on the top500.org list in terms of the number of sites and that's a result of that strategy.

We haven't gone as far as some of the PC vendors who are really just distributing PC Intel systems. My customers say to me that they would prefer to buy, our academics rather would prefer to buy the Intel systems from us because when they ring us up and they want to talk HPC and they want to talk MBI and grids etc., then people in Sun know what they're talking about. But a traditional PC company doesn't have that level of knowledge and that clearly drives a preference amongst the academics and they would like to talk to someone who has a HPC specialty

Con: So I guess the value in this fray is the intelligence invested in the software stack rather than the underlying complexity of the hardware?

Justin: Yes I agree with that, and that's our interest in this arrangement with University of Queensland and this installation, and that is that they are going to be using our software stack. Our value in this particular case is the grid engine as the technical compute portal and the control systems and obviously the Intel equipment. We're acquiring the CPU, were making the systems ourselves and they are going to have a unique configuration.

Con: And obviously overall with something like this there is the end to end integration of something of non-trivial complexity such as this.

Justin: Yes that's right!

Con: My reading of the market particularly in the past five years indicates what you guys are doing is very much where the market is going and the end point, I think if you go back five or seven years ago, the number of these kinds of clusters on top500.org and this is way back when it was starting up on the web, was very minimal, but nowadays it seems to be the case that eight out of ten new deployments are clusters, not because they are perfect for all situations but because of the value delivery, you get 'bang for the buck'. We seem to have overcome that mentality hump about requiring liquid Freon cooled big iron.

Justin: Yeh!

Con: Obviously any dis-advantages of this kind of architecture/platform OS and so-forth would have been far away beaten by the advantages, and that's why the customers are mandating this kind of approach, but from your perspective what are the dis-advantages of this kind of architecture at a technical level?

Justin: Well I know Steve would like to respond to this and I'll make one response.

We installed a couple of years ago, a similar cluster into University of Queensland mathematics department and that was a hundred and twenty eight single CPU machines and they were SPARC's CPU's with a hundred and twenty eight MB of memory on each one and this time round each node which has two CPUs with three GB of memory. And the customer said that when we interviewed them after we installed the equipment, how could we improve on what we delivered, and they said well the amount of memory per node increased would make it more usable and so we certainly done that this time round.

Although three GB memory on a node doesn't solve a certain class of problems and there are problems where customers would want to access twenty GB or a hundred GB of memory with a single problem or a single application and this is where the cluster approach a low-wide cluster is not well suited. So any large memory problems are likely be solved in a different way.

Con: Right so that's where you would be looking at more of a NUMA -type class machine?

Justin: Yes that's right or an SMP class

Con: Right, a complete shared memory system

Justin: We would go for our F15 K style product which would go to a hundred and six CPUs, we would use that could access a huge amount of memory.

And that is clearly solved within the scope of this particular proposal.

Con: Understood

Steve: Yeah I have to agree with Justin's assesment there Con, in terms of the large memory problems.

Con: So that would have been a break point if for this particular organisation this particular set of researchers if they had problems that needed very large in memory stores, very large files, that would have obviously been a problem with regards to the hardware architecture and for the interconnect architecture as well it just wouldn't have suited the kind of problem.

Obviously it wasn't an issue that's why they work with this kind of solution.

Steve: That is correct.

Con: How long does this kind of solution or rollout take? Obviously it is reasonably large, it's not huge in the extreme end and it's also not small, it's not just a couple of dozen boxes. Forgetting about all the business end of the equation, i.e the bidding phase and all that kinda stuff, but primarily from a technical perspective what kind of lapse time is required to get one of these babies going?

Justin: Well it's Justin here I can respond to that.

We would probably allow about a month to have the system commissioned, the reason why I'm hesitating is that there are some environmental things that need to be sorted out by the University so what that really means is that there is quite a deal of air-conditioning, inside preparations required and they're working through those issues at the moment. And frankly it's the environmental considerations that would potentially take longer than the actual implementation of the System.

Con: Wow, fast deployment!

Justin: Really once you have one image built for one node for the operating system then you can duplicate that fairly quickly across all of the nodes and physically it would take us a few days with a couple of engineers to unpack all the boxes and set them all up and then we would work through configuring the grid engine software and the portal software. So a month would be a reasonable allowance for that kind of project.

Con: I imagine that's actually one of the unbeknownst advantages of this kind of deployment, unlike a large vector processor system which will take a lot of gently-gently installation and configuration work, this kind of system seems to be pretty rock and ready to go very quickly. One of my friends was involved in the Cray deployment at the University of Queensland many years back. If I recall, he spent months there helping the team roll that system out.

Justin: Well I know about that Cray installation at the University of Queensland and there was a quite a bit of time it was required to get that system going because there where a lot of new things for the University to manage, it was water cooled, it was a new operating system for them, it was new CPUs, it was new compilers, there where new queuing systems. I mean everything was new, the operator console was new, you think about it.

With this current situation, really there is not a lot that's new to the University. They are already using Grid Engine, they're already signed up for our Sun ONE portal product which includes the technical compute portal and they're using portal, clearly they're using Linux, they are already using Xeon CPUs throughout the University and doing various jobs and Cisco switches which is the switch architecture. And really in terms of the hardware that's something that we're well trained at first step so we can set it up fairly quickly. They have also by the way implemented a 128-node system at mathematics as well so they quite a few good schools on-site so I think this will come up to speed fairly quickly.

Con: As you were mentioning that rather than going with the possibly un-necessarily expensive specialised hardware, this is very much into the realm of commoditisation of the underlying platform technology and their familiarity with the existing software, which is obviously making your job a lot easier for deploying these kinds of solutions. I suspect it's probably also making it far more affordable and increasingly likely that more and more organisations will pop-up and deploy such systems because .

At a technical level the grid engine technology you deploy, how inter-operable is that? How feasible is it to glue these disparate chunks of technology like grid engine and Globus together. As they are often built atop Java, so there is RMI possibilities and there is a possibility of SOAP as well, one would presume this is possible. Therefore, how inter-operable is this new installation to be glued into other systems than APAC?

Justin: In as much as you mentioned the standards, things like SOAP and XML-RPC etc., the inter-operability is there.

Justin: This has been successful for us in terms of where we got to today and the interest we've received from the University of Queensland and their willingness to work with us and we think that most of the other Universities in Australia will be very keen to have access the same kind of technology and so we are putting together a proposal which is going to create a grid throughout Australia using the grid engine product and you will see more about that in the future but that is likely to roll-out to other Universities over the next six months or so.

We are quite keen and we have been getting requests from other Universities to implement similar kinds of technology, maybe not at the same scale as a smaller scale but we are confident we could create a large grid or these sorts of systems.







Con Guys, thanks for your time today. Best of luck with these technologies in the future






Sun Thanks
team




AUUG 2003-The Photos

Anthony Rumble <smilie@LinuxHelp.com.au>









AUUG Chapter Meetings and Contact Details

CITY	LOCATION	OTHER
ADELAIDE	We meet at Internode, Level 3/132 Grenfell St aka 'the old AAMI building', at 7 pm on the second Wednesday of each month.	Contact sa-exec@auug.org.au for further details.
BRISBANE	Inn on the Park 507 Coronation Drive Toowong	For further information, contact the QAUUG Executive Committee via email (qauug-exec@auug.org.au). The 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.
CANBERRA	Australian National University	
HOBART	University of Tasmania	
MELBOURNE	Various. For updated information See: http://www.vic.auug.org.au/	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.
PERTH	The Victoria League 276 Onslow Road Shenton Park	
SYDNEY	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: http://www.auug.org.au/nswauug/

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	\$110.00 (including \$10 GST)	<input type="checkbox"/>
Renew/New* Student Membership	\$27.50 (including \$2.50 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:

This form serves as Tax Invoice. AUUG ABN 15 645 981 718